

MySQL Programmer's Guide

Ulrik Petersen

July 29, 2006

Copyright (C) 2001-2005 Ulrik Petersen

This document is made available under the Creative Commons Attribution-Sharealike license version 2.5.

See

<http://creativecommons.org/licenses/by-sa/2.5/>

for what that means.

Please visit the Emdros website for the latest news and downloads:

<http://emdros.org>

Abstract

This is the MQL Programmer's Guide. It documents Emdros version 1.2.0.pre172 and upwards. If you just wish to use Emdros to query your data, then this might not be for you. Instead, you can consult the MQL Query Guide, which is available from the Emdros website, or with any recent distribution of Emdros in the 1.2.0.preXX series.

In Chapter 1, we discuss some preliminaries, such as the history of Emdros, as well as giving an overview of the formalism used to define the MQL language, called Backus-Naur Form (or BNF).

In Chapter 2, we give an overview of the EMdF model, from a user's standpoint.

In Chapter 3, we describe the bulk of the MQL language.

In Chapter 4, we describe the query-subset of MQL. That is, the subset of MQL in which you can express queries that look for objects by object type, features, and structural relationships such as embedding and sequence.

Contents

1 Preliminaries	10
1.1 Introduction	10
1.2 Origins of MdF, EMdF, and MQL	10
1.3 Introduction to Backus-Naur Form	11
1.3.1 Context-Free Grammars	11
1.3.1.1 Rule	11
1.3.1.2 Non-terminal	12
1.3.1.3 Terminal	12
1.3.1.4 Choice	12
1.3.1.5 Concatenation	12
1.3.1.6 Start-symbol	12
1.3.2 Context-free grammars: Putting it all together	12
1.3.3 BNF	13
1.3.3.1 Introduction	13
1.3.3.2 Example	14
1.3.3.3 Elements of “MQL BNF”	14
1.4 Acknowledgements	15
1.4.1 Westminster Hebrew Institute	15
1.4.2 Regular expression support	15
2 The EMdF model	16
2.1 Introduction	16
2.2 Monads	16
2.2.1 Monads and textual sequence	16
2.2.2 Granularity	16
2.2.3 Text flow	17
2.2.4 Conclusion	17
2.3 Objects	17
2.3.1 What is an object?	17
2.4 Object types	17
2.4.1 Objects are grouped in types	17
2.5 Features	18
2.5.1 What is a feature?	18
2.5.2 Object types have features	18
2.5.3 Features have types	18
2.6 Example	19
2.7 Other concepts	19
2.7.1 Introduction	19
2.7.2 pow_m	20
2.7.3 any_m	20
2.7.4 all_m	20
2.7.5 object ids (id_d, id_m)	20

2.7.6	self	21
2.7.7	part_of	21
2.7.8	gaps	21
2.7.9	borders, first, last	21
2.7.10	consecutive with respect to a set of monads	21
2.7.11	enumerations	21
2.7.11.1	Definition	21
2.7.11.2	Example	21
2.7.11.3	Default constant	22
2.7.11.4	Terminology	22
2.7.11.5	Names are identifiers	22
2.7.11.6	Each enumeration is a name-space	22
2.7.11.7	Enumeration constants must be unique	22
2.7.12	min_m, max_m	22
2.7.13	Arbitrary monad sets	23
2.7.14	Databases	23
2.8	Encryption	23
3	MQL database manipulation	24
3.1	Preliminaries	24
3.1.1	Introduction	24
3.1.2	Non-terminals	24
3.1.3	Lexical conventions	24
3.1.4	Name-spaces	25
3.1.5	Top-level constraints on MQL syntax	26
3.2	Return types	26
3.2.1	Introduction	26
3.2.2	Output-formats	26
3.2.3	Tables	27
3.2.4	Atomic output-types in tables	27
3.2.5	Other return values	27
3.3	Database manipulation	28
3.3.1	CREATE DATABASE	28
3.3.1.1	Syntax	28
3.3.1.2	Example	28
3.3.1.3	Explanation	28
3.3.1.4	Return type	28
3.3.2	INITIALIZE DATABASE	28
3.3.2.1	Syntax	28
3.3.2.2	Example	29
3.3.2.3	Explanation	29
3.3.2.4	Return type	29
3.3.3	USE DATABASE	29
3.3.3.1	Syntax	29
3.3.3.2	Example	29
3.3.3.3	Explanation	30
3.3.3.4	Return type	30
3.3.4	DROP DATABASE	30
3.3.4.1	Syntax	30
3.3.4.2	Example	30
3.3.4.3	Explanation	30
3.3.4.4	Return type	31
3.3.5	VACUUM DATABASE	31
3.3.5.1	Syntax	31

3.3.5.2	Example	31
3.3.5.3	Explanation	31
3.3.5.4	Return type	31
3.3.6	DROP INDEXES	31
3.3.6.1	Syntax	31
3.3.6.2	Example	31
3.3.6.3	Explanation	32
3.3.6.4	Return type	32
3.3.7	CREATE INDEXES	32
3.3.7.1	Syntax	32
3.3.7.2	Example	32
3.3.7.3	Explanation	32
3.3.7.4	Return type	33
3.4	Transactions	33
3.4.1	BEGIN TRANSACTION	33
3.4.1.1	Syntax	33
3.4.1.2	Example	33
3.4.1.3	Explanation	33
3.4.1.4	Return type	33
3.4.2	COMMIT TRANSACTION	34
3.4.2.1	Syntax	34
3.4.2.2	Example	34
3.4.2.3	Explanation	34
3.4.2.4	Return type	34
3.4.3	ABORT TRANSACTION	34
3.4.3.1	Syntax	34
3.4.3.2	Example	34
3.4.3.3	Explanation	34
3.4.3.4	Return type	35
3.5	Object type manipulation	35
3.5.1	CREATE OBJECT TYPE	35
3.5.1.1	Syntax	35
3.5.1.2	Examples	36
3.5.1.3	Explanation	36
3.5.1.4	Return type	38
3.5.2	UPDATE OBJECT TYPE	38
3.5.2.1	Syntax	38
3.5.2.2	References	38
3.5.2.3	Example	39
3.5.2.4	Explanation	39
3.5.2.5	Return type	39
3.5.3	DROP OBJECT TYPE	39
3.5.3.1	Syntax	39
3.5.3.2	Example	39
3.5.3.3	Explanation	39
3.5.3.4	Return type	39
3.6	Enumeration manipulation	40
3.6.1	CREATE ENUMERATION	40
3.6.1.1	Syntax	40
3.6.1.2	References	40
3.6.1.3	Example	40
3.6.1.4	Explanation	40
3.6.1.5	Return type	41
3.6.2	UPDATE ENUMERATION	41

3.6.2.1	Syntax	41
3.6.2.2	References	41
3.6.2.3	Example	41
3.6.2.4	Explanation	42
3.6.2.5	Return type	42
3.6.3	DROP ENUMERATION	42
3.6.3.1	Syntax	42
3.6.3.2	Example	42
3.6.3.3	Explanation	42
3.6.3.4	Return type	42
3.7	Segment manipulation	42
3.7.1	Introduction	42
3.7.2	CREATE SEGMENT	43
3.7.2.1	Syntax	43
3.7.2.2	Example	43
3.7.2.3	Explanation	43
3.7.2.4	Return type	43
3.8	Querying	43
3.8.1	SELECT (FOCUS ALL) OBJECTS	43
3.8.1.1	Syntax	43
3.8.1.2	References	44
3.8.1.3	Example	44
3.8.1.4	Explanation	45
3.8.1.5	Monad set	45
3.8.1.6	Return type	45
3.8.2	SELECT OBJECTS AT	46
3.8.2.1	Syntax	46
3.8.2.2	Example	46
3.8.2.3	Explanation	46
3.8.2.4	Return type	46
3.8.3	SELECT OBJECTS HAVING MONADS IN	46
3.8.3.1	Syntax	46
3.8.3.2	References	46
3.8.3.3	Example	47
3.8.3.4	Explanation	47
3.8.3.5	Return type	47
3.8.4	GET OBJECTS HAVING MONADS IN	47
3.8.4.1	Syntax	47
3.8.4.2	References	48
3.8.4.3	Example	48
3.8.4.4	Explanation	48
3.8.4.5	Return type	48
3.8.5	SELECT OBJECT TYPES	48
3.8.5.1	Syntax	48
3.8.5.2	Example	49
3.8.5.3	Explanation	49
3.8.5.4	Return type	49
3.8.6	SELECT FEATURES	49
3.8.6.1	Syntax	49
3.8.6.2	Example	49
3.8.6.3	Explanation	49
3.8.6.4	Return type	49
3.8.7	SELECT ENUMERATIONS	50
3.8.7.1	Syntax	50

3.8.7.2	Example	50
3.8.7.3	Explanation	50
3.8.7.4	Return type	50
3.8.8	SELECT ENUMERATION CONSTANTS	50
3.8.8.1	Syntax	50
3.8.8.2	Example	50
3.8.8.3	Explanation	51
3.8.8.4	Return type	51
3.8.9	SELECT OBJECT TYPES USING ENUMERATION	51
3.8.9.1	Syntax	51
3.8.9.2	Example	51
3.8.9.3	Explanation	51
3.8.9.4	Return type	51
3.8.10	SELECT MIN_M	51
3.8.10.1	Syntax	51
3.8.10.2	Example	52
3.8.10.3	Explanation	52
3.8.10.4	Return type	52
3.8.11	SELECT MAX_M	52
3.8.11.1	Syntax	52
3.8.11.2	Example	52
3.8.11.3	Explanation	52
3.8.11.4	Return type	52
3.8.12	SELECT MONAD SETS	52
3.8.12.1	Syntax	52
3.8.12.2	Example	52
3.8.12.3	Explanation	52
3.8.12.4	Return type	53
3.8.13	GET MONAD SETS	53
3.8.13.1	Syntax	53
3.8.13.2	Example	53
3.8.13.3	Explanation	53
3.8.13.4	Return type	53
3.8.14	GET MONADS	53
3.8.14.1	Syntax	53
3.8.14.2	References	54
3.8.14.3	Example	54
3.8.14.4	Explanation	54
3.8.14.5	Return type	54
3.8.15	GET FEATURES	54
3.8.15.1	Syntax	54
3.8.15.2	References	55
3.8.15.3	Example	55
3.8.15.4	Explanation	55
3.8.15.5	Return type	55
3.8.16	GET SET FROM FEATURE	55
3.8.16.1	Syntax	55
3.8.16.2	Example	56
3.8.16.3	Explanation	56
3.8.16.4	Return type	56
3.9	Object manipulation	56
3.9.1	CREATE OBJECT FROM MONADS	56
3.9.1.1	Syntax	56
3.9.1.2	References	57

3.9.1.3	Example	57
3.9.1.4	Explanation	57
3.9.1.5	Return type	58
3.9.2	CREATE OBJECT FROM ID_DS	58
3.9.2.1	Syntax	58
3.9.2.2	References	58
3.9.2.3	Example	59
3.9.2.4	Explanation	59
3.9.2.5	Return type	59
3.9.3	CREATE OBJECTS WITH OBJECT TYPE	59
3.9.3.1	Syntax	59
3.9.3.2	References	59
3.9.3.3	Example	59
3.9.3.4	Explanation	60
3.9.3.5	Return type	60
3.9.4	UPDATE OBJECTS BY MONADS	61
3.9.4.1	Syntax	61
3.9.4.2	References	61
3.9.4.3	Example	61
3.9.4.4	Explanation	61
3.9.4.5	Return type	61
3.9.5	UPDATE OBJECTS BY ID_DS	62
3.9.5.1	Syntax	62
3.9.5.2	References	62
3.9.5.3	Example	62
3.9.5.4	Explanation	62
3.9.5.5	Return type	62
3.9.6	DELETE OBJECTS BY MONADS	62
3.9.6.1	Syntax	62
3.9.6.2	References	63
3.9.6.3	Example	63
3.9.6.4	Explanation	63
3.9.6.5	Return type	63
3.9.7	DELETE OBJECTS BY ID_DS	63
3.9.7.1	Syntax	63
3.9.7.2	References	63
3.9.7.3	Example	63
3.9.7.4	Explanation	63
3.9.7.5	Return type	64
3.10	Monad manipulation	64
3.10.1	MONAD SET CALCULATION	64
3.10.1.1	Syntax	64
3.10.1.2	References	64
3.10.1.3	Example	64
3.10.1.4	Explanation	65
3.10.1.5	Return type	65
3.10.2	CREATE MONAD SET	65
3.10.2.1	Syntax	65
3.10.2.2	References	65
3.10.2.3	Example	65
3.10.2.4	Explanation	65
3.10.2.5	Return type	65
3.10.3	UPDATE MONAD SET	66
3.10.3.1	Syntax	66

3.10.3.2	References	66
3.10.3.3	Examples	66
3.10.3.4	Explanation	66
3.10.3.5	Return type	67
3.10.4	DROP MONAD SET	67
3.10.4.1	Syntax	67
3.10.4.2	Example	67
3.10.4.3	Explanation	67
3.10.4.4	Return type	67
3.11	Meta-statements	67
3.11.1	QUIT	67
3.11.1.1	Syntax	67
3.11.1.2	Example	67
3.11.1.3	Explanation	67
3.11.1.4	Return type	67
4	MQL Query subset	68
4.1	Introduction	68
4.2	Informal introduction to MQL by means of some examples	68
4.2.1	Introduction	68
4.2.2	topograph	69
4.2.3	features	69
4.2.4	object_block, object_block_first	69
4.2.5	power	70
4.2.6	opt_gap_block	71
4.2.7	gap_block	71
4.2.8	object references	72
4.3	Syntax of mql_query	72
4.3.1	Introduction	72
4.3.2	Syntax	72
4.3.3	References	74
4.4	The sheaf	75
4.4.1	Introduction	75
4.4.2	Structure of the sheaf	75
4.4.2.1	What is a sheaf?	75
4.4.2.2	What is a straw?	75
4.4.2.3	What is a matched_object?	75
4.4.3	MQL is topographic	75
4.4.4	Meaning of matched_object	76
4.4.5	Meaning of straw	76
4.4.6	Meaning of the sheaf	76
4.4.7	Flat sheaf	77
4.5	Universe and substrate	77
4.5.1	Introduction	77
4.5.2	Universe and substrate	77
4.6	Consecutiveness and embedding	78
4.7	Blocks	78
4.7.1	Introduction	78
4.7.2	Object blocks	78
4.7.2.1	Introduction	78
4.7.2.2	Syntax	79
4.7.2.3	References	79
4.7.2.4	Examples	79
4.7.2.5	Explanation	80

4.7.3	Gap blocks	81
4.7.3.1	Introduction	81
4.7.3.2	Syntax	81
4.7.3.3	Examples	81
4.7.3.4	Explanation	82
4.7.4	Retrieval	82
4.7.4.1	Introduction	82
4.7.4.2	Syntax	82
4.7.4.3	Examples	82
4.7.4.4	Explanation	83
4.7.5	First and last	83
4.7.5.1	Introduction	83
4.7.5.2	Syntax	83
4.7.5.3	Examples	84
4.7.5.4	Explanation	84
4.7.6	Feature constraints	84
4.7.6.1	Introduction	84
4.7.6.2	Syntax	84
4.7.6.3	References	85
4.7.6.4	Examples	85
4.7.6.5	Explanation of Examples	86
4.7.6.6	Explanation	86
4.7.6.7	Type-compatibility	87
4.7.6.8	Comparison-operators	87
4.7.7	Object references	89
4.7.7.1	Introduction	89
4.7.7.2	Syntax	89
4.7.7.3	Examples	89
4.7.7.4	Explanation of examples	89
4.7.7.5	Explanation	90
4.7.7.6	Constraints on object references	90
4.7.8	Block	91
4.7.8.1	Introduction	91
4.7.8.2	Syntax	91
4.7.8.3	The “*” construct	91
4.8	Strings of blocks	92
4.8.1	Introduction	92
4.8.2	topograph	93
4.8.2.1	Introduction	93
4.8.2.2	Syntax	93
4.8.2.3	Examples	93
4.8.2.4	Explanation of examples	93
4.8.2.5	Universe and Substrate	94
4.8.3	blocks	94
4.8.3.1	Introduction	94
4.8.3.2	Syntax	94
4.8.4	block_string	94
4.8.4.1	Introduction	94
4.8.4.2	Syntax	94
4.8.4.3	Examples	94
4.8.4.4	Explanation	95
4.8.5	block_str	95
4.8.5.1	Introduction	95
4.8.5.2	Syntax	95

4.8.5.3	Examples	95
4.8.5.4	power	96
4.8.5.5	The bang (“!”)	96
A	Copying	97
A.1	Introduction	97
A.2	Creative Commons Deed (for all documentation)	97
A.3	GNU General Public License	101
A.4	PCRE license	106
A.4.1	THE BASIC LIBRARY FUNCTIONS	106
A.4.2	THE C++ WRAPPER FUNCTIONS	107
A.4.3	The “BSD” license	107
B	Console sheaf grammar	108
B.1	Introduction	108
B.2	Sheaf grammar	108
B.3	References	109

Chapter 1

Preliminaries

1.1 Introduction

Welcome to the MQL Programmer’s Guide. MQL is the query language associated with the EMdF model. MQL is a “full access language,” which simply means that MQL lets you create, update, delete, and query most of the data domains in the EMdF model – databases, objects, object types, features, etc. MQL is your front-end to the EMdF database engine. This guide helps you formulate MQL queries that do what you need done with your EMdF database.

This guide has four chapters. The first is this chapter on preliminaries. The second is a gentle introduction to the EMdF model, which underlies the MQL language. The third chapter deals with the bulk of the MQL query language, detailing all the different kinds of queries for creating, updating, deleting, and querying an EMdF database. The fourth chapter is a special chapter devoted to explaining those MQL queries that query for objects in the database. Since these queries are so rich, it was deemed necessary to devote a whole chapter to their treatment.

This chapter will proceed as follows. First, we present a short history of MdF, EMdF, and MQL. Second, we give a gentle introduction to Backus-Naur Form, or BNF, which will be used throughout chapters 3 and 4. Lastly, we explain the origin of the support for something called *regular expressions* in MQL. This is done so as to comply with the license for the library used.

But first, an explanation of where EMdF and MQL come from.

1.2 Origins of MdF, EMdF, and MQL

EMdF and MQL are not original works. They are merely derivative works based on someone else’s hard labors. Most of the ideas underlying the EMdF model and the MQL query language are to be found in the PhD thesis of Crist-Jan Doedens, published in 1994 as [Doedens94]. This thesis described, among other things, the MdF database model and the QL query language. As one might guess, EMdF stems from MdF, and MQL stems from QL.

The EMdF model takes over the MdF model in its entirety, but adds a few concepts which are useful when implementing the MdF model in real life. Thus the ‘E’ in ‘EMdF’ stands for “Extended”, yielding the “Extended MdF model”. The EMdF model is the subject of chapter 2.

“MQL” stands for “Mini QL.” Originally, I devised MQL as a subset of the QL query language developed in Dr. Doedens’ PhD thesis, hence the “Mini” modifier. Since then, however, MQL has grown. QL was not a full access language, but specified only how to query an MdF database, i.e., how to ask questions of it. MQL, by contrast, is a full access language, allowing not only querying, but also creation, update, and deletion of the data domains of the EMdF model. The MQL query language is the subject of chapters 3 and 4.

Thus EMdF and MQL are derivatives of the MdF database model and the QL query language developed by Dr. Crist-Jan Doedens in his 1994 PhD thesis.

1.3 Introduction to Backus-Naur Form

1.3.1 Context-Free Grammars

BNF is a way of specifying the “syntactic rules” of a language. English also has “syntactic rules,” and some of them can be specified using a “Context-Free Grammar.” BNF is precisely a way of specifying a context-free grammar for a formal language. Thus it is beneficial first to see what a context-free grammar is, before looking at the details of BNF.

In English, the basic clause-pattern is “Subject - Verb - Object”. For example, in the clause “I eat vegetables,” the word “I” is the subject, the word “eat” is the verb, and the word “vegetables” is the object. A clause which exhibits exactly the same “Subject - Verb - Object” structure is “You drink coke.” Here, “You” is the subject, “drink” is the verb, and “coke” is the object.

Consider the following context-free grammar:

```
Sentence → NPsubj VP
NPsubj → "I" | "You"
VP → V NPobj
V → "eat" | "drink"
NPobj → "vegetables" | "coke"
```

This little context-free grammar is a toy example of a context-free grammar. However, despite its simplicity, it exemplifies all of the concepts involved in context-free grammars:

- Rule
- Non-terminal
- Terminal
- Choice
- Concatenation
- Start-symbol.

These will be described in turn below

1.3.1.1 Rule

A “Rule” consists of three parts:

1. The left-hand side
2. The “production arrow” (“→”).
3. The right-hand side

An example of a rule in the above context-free grammar is:

```
Sentence → NPsubj VP
```

It specifies that the left-hand side (“Sentence”) can be *replaced with* the right-hand side, which in this case is two symbols, “NP_{subj}” followed by “VP”. Sometimes, we also say that a left-hand side is *expanded to* the right-hand side.

1.3.1.2 Non-terminal

There are only two kinds of symbols in a context-free grammar: Non-terminals and Terminals. They are a contrasting pair. In this section, we describe what a non-terminal is, and in the next section, what a terminal is.

A “Non-terminal” is a symbol in a rule which can be expanded to other symbols. Thus the symbols “Sentence”, “NP_{subj}”, “VP”, “V”, and “NP_{obj}” constitute all of the non-terminals of the above context-free grammar.

Only non-terminals can stand on the left-hand side of a rule. A non-terminal is defined as a symbol which can be expanded to or replaced with other symbols, and hence they can stand on the left-hand side of a rule. But as you will notice in the above context-free grammar, a non-terminal can also stand on the right-hand-side of a rule. For example, the non-terminal “V” is present both in the rule for how to expand the non-terminal “VP”, and in the rule for how to expand itself. Thus, in order to expand “VP” fully, you must first expand to the right-hand-side “V NP_{obj}”, and then expand both “V” and “NP_{obj}”, using the rules for these two.

1.3.1.3 Terminal

A “Terminal” is a symbol in a rule which cannot be expanded to other symbols. Hence, it is “terminal” in the sense that the expansion cannot proceed further from this symbol. In the above context-free grammar, the terminals are: “I”, “You”, “eat”, “drink”, “vegetables”, and “coke”. These are symbols which cannot be expanded further.

Terminals can only stand on the right-hand side of a rule. If they were to stand on the left-hand-side of the rule, that would mean that they could be expanded to or replaced with other symbols. But that would make them non-terminals.

1.3.1.4 Choice

In the rule for “V” in the above grammar, we see an example of choice. The choice is indicated by the “|” symbol, which is read as “or”. Thus, this example:

$$V \rightarrow \text{"eat"} \mid \text{"drink"}$$

is read as ‘V expands to “eat” or “drink”’.

1.3.1.5 Concatenation

We have already seen an example of concatenation, namely in the rule for “Sentence”:

$$\text{Sentence} \rightarrow \text{NP}_{\text{subj}} \text{VP}$$

Here, the symbols “NP_{subj}” and “VP” are *concatenated*, or placed in sequence. Thus “VP” comes immediately after “NP_{subj}”.

“Concatenated” is simply a fanciful name for “being in sequence”, but although it is a basic idea, we included it for completeness.

1.3.1.6 Start-symbol

The start-symbol is merely the left-hand side non-terminal of the first rule in the grammar. Thus, in the above grammar, “Sentence” is the start-symbol.

1.3.2 Context-free grammars: Putting it all together

It is time to see how all of this theory works in practice. The above grammar can produce 8 sentences, some of which do not make sense:

1. I eat vegetables

2. I eat coke
3. I drink vegetables
4. I drink coke
5. You eat vegetables
6. You eat coke
7. You drink vegetables
8. You drink coke

Let us pick one of these sentences and see how it was produced from the above grammar. We will pick number 8, “You drink coke”, and trace all the steps. We start with the start-symbol, “Sentence”:

1. Sentence
This is expanded using the rule for “Sentence”:
2. NP_{subj} VP
The “ NP_{subj} ” non-terminal is then expanded to “You” using one of the choices in the rule for NP_{subj} :
3. “You” VP
The “VP” is then expanded using the rule for “VP”:
4. “You” V NP_{obj}
The “V” non-terminal is then expanded to the terminal “drink”:
5. “You” “drink” NP_{obj}
The “ NP_{obj} ” non-terminal is then expanded to the terminal “coke”:
6. “You” “drink” “coke”
Which yields the final sentence, “You drink coke”. This sentence has no non-terminals, only terminals, and therefore it cannot be expanded further. We are finished.

If you would like to, try to trace the production of one of the other sentences using pencil and paper, tracing each step as in the above example. When you have done so once or twice, you should understand all there is to understand about context-free grammars.

And BNF is simply a way of specifying a context-free grammar. So let us start looking at the details of BNF.

1.3.3 BNF

1.3.3.1 Introduction

BNF comes in various variants, and almost everyone defines their usage of BNF a little differently from everyone else. In this document, we shall also deviate slightly from “standard BNF”, but these deviations will only be very slight.

This treatment of BNF will be made from an example of a context-free grammar in “MQL BNF.” This example covers every formalism used in “MQL BNF,” and is a real-life example of the syntax of an actual MQL statement:

1.3.3.2 Example

```
create_enumeration_statement : "CREATE"  
    ("ENUMERATION" | "ENUM")  
    enumeration_name "="  
    "{" ec_declaration_list "  
;  
enumeration_name : T_IDENTIFIER  
    /* The T_IDENTIFIER is a terminal  
    denoting an identifier. */  
;  
ec_declaration_list : ec_declaration { "," ec_declaration }  
;  
ec_declaration : [ "DEFAULT" ]  
    ec_name [ ec_initialization ]  
;  
ec_name : T_IDENTIFIER  
;  
ec_initialization : "=" T_INTEGER  
;
```

1.3.3.3 Elements of “MQL BNF”

All of the elements of “MQL BNF” can be listed as follows:

1. Rule

This is just the same as in the context-free grammars above, except that the “→” production arrow is replaced by a “:” colon. Also, a rule ends with a “;” semicolon.

2. Non-terminal

Non-terminals always start with a lower-case letter, e.g., “ec_declaration.”

3. Terminal

Terminals are either strings in “double quotes” or strings that start with “T_”, e.g., “T_IDENTIFIER”.

4. Choice

This is the same as in the context-free grammars. The symbol “|” is used.

5. Concatenation

This is the same as in the context-free grammars. A space between two symbols is used.

6. Start-symbol

This is the same as in the context-free grammars. The left-hand side non-terminal of the first rule is the start-symbol.

7. Comment

A comment is enclosed in /* slashes and stars */. The comments are not part of the grammar, but serve to explain a part of the grammar to the human reader.

8. Optional specification

A symbol or sequence of symbols enclosed in [square brackets] is considered optional. For example, the ec_initialization non-terminal is optional in the ec_declaration rule. Both terminals and non-terminals can be optional. Here, the choice is between “ENUM” and “ENUMERATION”, rather than, say, “CREATE ENUMERATION” and the rest of the rule.

9. Bracketing

Sometimes, we do not go to the trouble of spelling out a choice with a rule by itself. Instead, (brackets) are used. For example, in the above grammar, there is a choice between “ENUMERATION” and “ENUM”. Only one must be chosen when writing the CREATE ENUMERATION statement. The brackets serve to let you know the scope of the choice, i.e., between exactly which elements the choice stands.

10. Bracing (repetition)

The { braces } are used to indicate that the enclosed part can be repeated zero or more times. For example, in the rule for `ec_declaration_list`, the first `ec_declaration` can be followed by zero or more occurrences of first a “,” comma and then an `ec_declaration`. This effectively means that the `ec_declaration_list` is a comma-separated list of `ec_declarations`, with one or more `ec_declarations`.

There is no comma after the last `ec_declaration`. To see why, notice that the part that is repeated starts with a comma and ends with an `ec_declaration`. Thus, no matter how many times you repeat this sequence, the `ec_declaration` will always be last, and the comma will never be last.

That it can be repeated *zero* or more times simply means that it is optional. In the rule for `ec_declaration_list`, the first `ec_declaration` can stand alone.

Notice also that, in the rule for `create_enumeration_statement`, there are braces as well. These braces, however, are enclosed in “double quotes” and are therefore terminals. Thus they do not have the meaning of repetition, but are to be written in the statement when writing a CREATE ENUMERATION statement. The braces without double quotes are repetition-braces and must not be written.

1.4 Acknowledgements

1.4.1 Westminster Hebrew Institute

Emdros development is made possible in part by the generosity of the Westminster Hebrew Institute¹, part of Westminster Theological Seminary², Philadelphia, Pennsylvania. The WHI grants me access to some seriously powerful hardware, which helps drive Emdros development, in particular the scalability of Emdros.

1.4.2 Regular expression support

MQL has support for regular expressions in queries. Regular expression support is provided by the PCRE library package, which is open source software, written by Philip Hazel, and copyright by the University of Cambridge, England. The PCRE library can be downloaded from:

`ftp.csx.cam.ac.uk/pub/software/programming/pcre/`

The PCRE included with Emdros is a modified copy.

We'll get back to regular expressions in section 4.7.6.8.4 and in section A.4. See also the index.

¹See <<http://whi.wts.edu/WHI>>.

²See <<http://www.wts.edu>>.

Chapter 2

The EMdF model

2.1 Introduction

This chapter is a gentle introduction to the “EMdF model”. The EMdF model is a “database model.” As such, it provides a solid theoretical foundation for the EMdF database engine and the MQL query language. However, its importance goes beyond being a theoretical foundation. The EMdF model defines the very way we talk about text in an EMdF database, and as such, it provides the “vocabulary” of the language by which you, the user, communicate with the EMdF database engine. Put another way, the EMdF model defines the concepts which you use to talk about text when communicating with the EMdF database engine. Thus it is very important that you understand all of the concepts involved in the EMdF model. However, these concepts are neither many nor difficult to understand. This chapter is designed to help you understand them, whatever your background.

This chapter is built around the concepts in the EMdF model: monads, objects, object types, features, and a few other concepts. These four concepts: monads, objects, object types, and features, form the backbone of the EMdF model. Once you have understood these, the rest are mere extensions which follow fairly easily.

And so, with no further ado, let us jump into the first of the four major concepts, monads.

2.2 Monads

2.2.1 Monads and textual sequence

Language is linear in nature. We produce language in real-time, with sequences of sounds forming sequences of words. Text, being language, is also linear in nature. In the EMdF model, this linearity is captured by the monads.

A monad is simply an integer – no more, no less. The sequence of integers (1,2,3,4,...) forms the backbone to which the flow of the text is tied. Thus, because a monad is an integer, and because we can order the integers in an unambiguous way, we use the sequence of monads to keep track of the sequence of textual information.

The sequence of monads begins at 1 and extends upwards to some large number, depending on the size of the database.

2.2.2 Granularity

What unit of text does a monad correspond to? For example, does a monad correspond to a morpheme, a word, a sentence, or a paragraph?

The answer is that you, the database user, decide the granularity of the EMdF database. You do this before any text is put into the database. If you want each monad to correspond to a morpheme, you simply decide that this is so. A more common choice is for each monad to correspond to a word. However, there

is nothing implicit or explicit in the EMdF model that prevents the user from deciding that another unit of text should correspond to a monad. Be aware, however, that once the choice has been made, and the database has been populated with text, it is not easy to revoke the decision, and go from, say, a word-level granularity to a morpheme-granularity.

2.2.3 Text flow

What is the reading-order of an EMdF database? Is it left-to-right, right-to-left, top-to-bottom, or something else?

The answer is that the EMdF model is agnostic with respect to reading-order. In the EMdF model, what matters is the textual sequence, as embodied by the monads. How the text is displayed on the screen is not specified in the EMdF model.

The MQL query language provides for both 7-bit (ASCII) and 8-bit encodings of strings, which means that the database implementor can implement any character-set that can be encoded in an 8-bit encoding, including Unicode UTF-8.

2.2.4 Conclusion

A monad is an integer. The sequence of integers (1,2,3,4,...) dictates the textual sequence. The granularity of an EMdF database is decided by the database-implementor. The EMdF model is agnostic with respect to reading-order (right-to-left, left-to-right, etc.).

2.3 Objects

2.3.1 What is an object?

An object is a set of monads. Thus, for example, an object might consist of the monads {1,2,4}. This object could, for example, be a phrase-object consisting of three words (assuming the monad-granularity is “one monad, one word”).

The EMdF model does not impose any restrictions on the set of monads making up an object. For example, objects can be discontinuous, as in the above example. In addition, objects can have exactly the same monads as other objects, and objects may share monads.

We need the last two concepts, object types and features, before we can understand exactly how an object can encode, say, a word or a phrase.

2.4 Object types

2.4.1 Objects are grouped in types

In any populated EMdF database, there will be at least one object type. Otherwise, no objects can be created, and thus the database cannot store textual information.

Objects are grouped in types, such as “word”, “phrase”, “clause”, “sentence”, but also “chapter”, “part”, “page”, etc. Each of these are potential object types that the user can create. Once an object type has been created, objects of that type can also be created.

Any object is of exactly one object type. Object types are what gives objects their interpretation. For example, an object of type “phrase” is, by itself, just a set of monads, such as {1,2,4}. But seen in conjunction with its object type, it becomes possible to interpret those monads as a number of words that make up a phrase.

An object type is also what determines what *features* an object has. And thus we turn to the last major concept in the EMdF model, namely features.

2.5 Features

2.5.1 What is a feature?

A feature is a way of assigning data to an object. For example, say we have an object of type “word”. Let us call this object “O”, and let us say that it consists of the singleton monad set {1}. Assume, furthermore, that the object type “word” has a feature called “surface_text”. Then this feature, taken on the object O, might be the string “In”. This is denoted as “O.surface_text”. If we have another object, O₂, which consists of the singleton monad set {2}, the value O₂.surface_text might be “the”. Thus we know that the first text in this EMdF database starts with the words “In the”.

2.5.2 Object types have features

An object type has a fixed number of features defined by the database implementor. For example, it might be necessary for a particular application to have these features on the object type “word”:

- surface_text
- lexical_form
- part_of_speech
- preceding_punctuation
- trailing_punctuation
- ancestor

The “ancestor” feature would be a pointer to another object, allowing the user to create an immediate constituency hierarchy.

The object type “phrase” might have the following list of features:

- phrase_type
- ancestor

2.5.3 Features have types

An object type has a number of features. Each feature, in turn, has one type. In the current implementation of the EMdF model, a feature can have one of the following types:

- integer
- string (which is an 8-bit string)
- ascii (which is a 7-bit ASCII string)
- id_d (which is an object id – we will get to this later)
- enumeration (which we will also get to later in this chapter)
- list of integer
- list of id_d
- list of enumeration constants

	1	2	3	4	5	6	7	8	9
Word	1	2	3	4	5	6	7	8	9
surface	The	door,	which	opened	towards	the	East,	was	blue.
part_of_speech	def.art.	noun	rel.pron.	verb	prep.	def.art.	noun	verb	adject.
Phrase	1		2	3		5		6	7
phrase_type	NP		NP	VP		NP		VP	AP
Phrase					4				
phrase_type					PP				
Clause_atom	1		2				3		
Clause	1		2				1		
Sentence	1								

Figure 2.1: Exemplifying the four major concepts of the EMdF model: monad, object, object type, and feature.

2.6 Example

We have now defined all of the four major concepts in the EMdF model: monad, object, object type, and feature. It is time to make them more concrete by giving an example of a very small EMdF database. Look at figure 2.1. At the top of the figure, you see the sequence of monads, 1 to 9. Just below this sequence, you see the object type “Word” with its object ordinals, 1 to 9. In this example, the granularity is “one monad, one word.” Thus *Word* number 1 corresponds to *monad* number 1, but they are really separate entities. *Word* number 1 *consists of* the set of monads {1}.

This becomes clearer when you notice that *Clause* number 2 consists of the set of monads {3,4,5,6,7}. Thus there is a fundamental distinction between the number of an object (also called object ordinal), and the set of monads making up that object.

Some of the object types in the figure (Word and Phrase) have a number of *features*. The object type “Word” has the features “surface” and “part_of_speech”. The “Phrase” object type has only one feature, namely “phrase_type”.

Notice that objects can be discontinuous. The *Clause* object with object ordinal 1 consists of the monads {1,2,8,9}. Thus there can be gaps in an object.

Notice also that an object type need not have features. The object types *Clause_atom*, *Clause*, and *Sentence* have no features in the figure.

2.7 Other concepts

2.7.1 Introduction

Having learned the basic concepts of the EMdF model, we now turn to the additional concepts which we use to talk about EMdF databases. These concepts are:

1. The special object types *pow_m*, *any_m*, and *all_m*
2. object ids (*id_d*, *id_m*)

3. self
4. part_of
5. gaps
6. borders, first, and last
7. enumerations
8. min_m and max_m
9. arbitrary monad sets
10. databases

2.7.2 pow_m

In each EMdF database, we assume an abstract object type, `pow_m`. This object type has one object for every possible set of monads. Thus the `pow_m` object type has objects consisting of the sets $\{1\}, \{2\}, \{3\}, \dots, \{1,2\}, \{1,3\}, \{1,4\}, \dots, \{2,3\}, \{2,4\}, \{2,5\}, \dots, \{1,2,3\}, \{1,2,4\}, \dots$, etc. Every possible set of monads is represented in the `pow_m` object type.

The `pow_m` object type is an *abstract* object type. That is, no objects of type `pow_m` actually exist in the EMdF database. However, it is useful to be able to talk about a particular `pow_m` object. In effect, a `pow_m` object is simply a set of monads, and sometimes, it is convenient to be able to talk about a particular `pow_m` object. This is especially true with gaps (see below).

The `pow_m` object type has no features.

2.7.3 any_m

The `any_m` object type is an abstract object type like `pow_m`. Each of its objects consist of a single monad. So the `any_m` objects are: $\{1\}, \{2\}, \{3\}, \dots$ etc. The `any_m` object type has no features.

2.7.4 all_m

The `all_m` object type has only one object, and it consists of all the monads in the database. That is, it consists of the monads from `min_m` to `max_m` (see sections 3.8.10 on page 51 and 3.8.11 on page 52), the smallest and the largest monads in use in the database at any given time. This one object is called `all_m-1`.

2.7.5 object ids (`id_d`, `id_m`)

Each object in the database (apart from `pow_m` objects) has an *object id_d*. An object `id_d` is a unique ID assigned to the object when the object is created. The `id_d` is used only for that particular object, and the `id_d` is never used again when the object is deleted.

A feature can have the type “`id_d`”, meaning that the values of the feature are taken from the `id_ds` in the database.

Each object in the database (including `pow_m`, `any_m`, and `all_m` objects) also has an `id_m`. The `id_m` is simply the set of monads which makes up the object. This is not strictly an ID, since objects of the same object type may have exactly the same monads. However, for historical reasons, this is called an `id_m`. See [Doedens94] or [Standard-MdF] for details.

2.7.6 self

Each object type in the database (apart from the `pow_m`, `any_m`, and `all_m` object types) has a feature called “self”. This is used to get the object `id_d` of the object in question.

The “self” feature is a read-only feature, which means that you cannot update an object’s self feature, or write to it when creating the object. The value of the “self” feature is assigned automatically when the object is created.

The type of the “self” feature is “`id_d`”.

2.7.7 part_of

If all of the monads of one object, O_1 , are contained within the set of monads making up another object, O_2 , we say that O_1 is `part_of` O_2 .

For example, an object with the monads $\{1,2\}$ would be `part_of` another object with the monads $\{1,2,4\}$.

In mathematical terms, O_1 is `part_of` O_2 if and only if $O_1 \subseteq O_2$.

2.7.8 gaps

Objects may have gaps. A gap in an object is a maximal stretch of monads which are not part of the object, but which are nevertheless within the boundaries of the endpoints of the object. For example, an object consisting of the monads $\{1,3,4,7,8,13\}$ has three gaps: $\{2\}$, $\{5,6\}$, and $\{9,10,11,12\}$.

Note that gaps are always maximal, i.e., extend across the whole of the gap in the object. For example, $\{6\}$ is not a gap in the above object: instead, $\{5,6\}$ is.

2.7.9 borders, first, last

Each non-empty object, being a set of monads, has a left border and a right border. The left border is the lowest monad in the set, while the right border is the highest monad in the set. These are also called the first monad and the last monad in the object. If we have an object O , the notation for these monads is `O.first` and `O.last`.

For example, if we have an object O consisting of the monads $\{2,3,4,5\}$, then `O.first` = 2 and `O.last` = 5.

2.7.10 consecutive with respect to a set of monads

The basic idea is that two sets of monads are consecutive if they follow each other without any gaps in between. However, this idea is extended so that the “no gaps in between” is interpreted with respect to a reference set of monads called S_u . For example, if $S_u = \{1,2,5,6\}$, then the sets $\{2\}$ and $\{5\}$ are consecutive with respect to S_u . However, the sets $\{2\}$ and $\{6\}$ are not consecutive with respect to S_u , since there is a “gap” consisting of the monad 5 in between the two sets. Likewise, the sets $\{1\}$ and $\{5\}$ are not consecutive with respect to S_u , because S_u has a monad, 2, which is a “gap” between the two sets.

2.7.11 enumerations

2.7.11.1 Definition

Each feature, it will be remembered, is of a certain type. These can be integers, strings, and `id_ds`, but they can also be enumerations. An enumeration is a set of pairs, where each pair consists of a constant-identifier and an integer value.

2.7.11.2 Example

For example, the enumeration “`phrase_type_t`” might have the pairs of constants and values as in table 2.1 on the following page.

constant	value
phrase_type_unknown	-1
VP	1
NP	2
AP	3
PP	4
AdvP	5
ParticleP	6

Table 2.1: phrase_type_t enumeration

2.7.11.3 Default constant

Each enumeration has exactly one default constant which is used when the user does not give a value for a feature with that enumeration type. In this example, “phrase_type_unknown” might be the default.

2.7.11.4 Terminology

The constants are called *enumeration constants*, while the type gathering the enumeration constants into one whole is called an *enumeration*.

2.7.11.5 Names are identifiers

The names of both enumerations and enumeration constants must be *identifiers*. See section 3.1.3 on page 24 for information on what an identifier is.

2.7.11.6 Each enumeration is a name-space

Each enumeration forms its own namespace. All name-spaces in the MQL language are orthogonal to each other. This means that two enumeration constants within the same enumeration cannot be called by the same constant-identifier, but two enumeration constants in two different enumerations may be the same. For more information, see section 3.1.4 on page 25 for more information.

2.7.11.7 Enumeration constants must be unique

Enumeration constants must be unique within each enumeration, both in their values and in their names. For example, you cannot have two labels with the same name in the same enumeration. Nor can you have two labels with the same value in the same enumeration, even if the labels have different names.

This is different from C or C++ enumerations, where the same value can be assigned to different labels.

Thus an enumeration is effectively a one-to-one correspondence (also called a bijective function) between a set of label names and a set of values.

2.7.12 min_m, max_m

An EMdF database has a knowledge of which is the smallest monad in use (min_m) and which is the largest monad in use (max_m). Normally, you don’t need to worry about these; the database maintains its knowledge of these monads without your intervention. You can, however, query the database for the minimum and maximum monads (see sections 3.8.10 on page 51 and 3.8.11 on page 52), and when you query the database for objects (section 3.8.1 on page 43), this is done within the confines of the minimum and maximum monads. Thus it is useful to know of their existence, but you needn’t worry too much about them.

The associated statements are SELECT MIN_M (section 3.8.10 on page 51) and SELECT MAX_M (section 3.8.11).

2.7.13 Arbitrary monad sets

Each database has a central repository of monad sets which are not associated with any objects. That is, they are not objects, have no object type, and no features. They are just plain monad sets.

These monad sets can be used as the basis for searches. That is, when doing a SELECT ALL OBJECTS query (or SELECT FOCUS OBJECTS), one can specify within which arbitrary monad set the search should be conducted.

The associated statements are SELECT MONAD SETS (section 3.8.12 on page 52), GET MONAD SETS (section 3.8.13 on page 53), CREATE MONAD SET (section 3.10.2 on page 65), UPDATE MONAD SET (section 3.10.3 on page 66), and DROP MONAD SET (section 3.10.4 on page 67).

2.7.14 Databases

The EMdF model has a concept of “database.” It is an organizational concept which generally corresponds to what the back-end database system calls a “database.” Within a database, there is one string of monads starting at 1 and extending upwards to some very large number. Within this stretch of monads, the user is free to create objects.

You may need to issue the USE DATABASE statement (see section 3.3.3 on page 29) as the first thing you do before doing anything else, in order to tell Emdros which database you want to deal with. Ask the implementor of your Emdros application whether this is what you should do.

A database can be created with the CREATE DATABASE statement (see section 3.3.1 on page 28).

2.8 Encryption

Dr. D. Richard Hipp, the author of SQLite, makes an encryption-enabled version available for a fee. There is skeleton support for SQLite encryption in Emdros, meaning one should be able to use Dr. Hipp’s encryption-enabled version of SQLite with Emdros and get encryption-support in Emdros on SQLite. This has not been tested, however; only the skeleton is there. See <<http://www.hwaci.com>>, the website of Dr. Hipp’s consulting company, for more information about Dr. Hipp’s encryption.

In this manual, when we speak of “encryption” on SQLite, please be aware that the actual encryption is not a part of Emdros, and you will achieve the exact same results and generate the exact same Emdros databases whether you use a key or not, unless you obtain an encryption-enabled SQLite from somewhere.

Chapter 3

MQL database manipulation

3.1 Preliminaries

3.1.1 Introduction

In this section on preliminaries, we will talk about four things. First, we describe what non-terminals are used in the grammar-fragments in this manual. Second, we describe the lexical conventions of MQL. Third, we describe the name-spaces available in MQL. And finally, we describe some top-level constraints in MQL syntax.

3.1.2 Non-terminals

The following non-terminals are used in this grammar:

- T_IDENTIFIER
- T_INTEGER
- T_STRING
- "strings", e.g., "OBJECT".

The fourth kind, e.g., "OBJECT" represent keywords in MQL. They are parsed as case-insensitive strings without the quotes.

3.1.3 Lexical conventions

The lexical conventions for MQL are as follows:

1. There are two kinds of comments:
 - (a) Enclosed in `"/*` (opening) and `*/` (closing). This kind of comment can span multiple lines. This is the same as C-style comments.
 - (b) Starting with `//` and ending at the end of the line. This is the style used in C++.
2. All keywords (such as `"CREATE"`, `"SELECT"`, `"<="`, etc.) are case-*insensitive* insofar as they are made up of letters. Keywords are enclosed in "double quotes" in the syntax-sections below.
3. A T_IDENTIFIER begins with any letter (a-z,A-Z) or an underscore (`_`), followed by zero or more letters (a-z,A-Z), numbers (0-9), or underscores (`_`). For example, `"Word"`, `"_delimiter"`, `"i18n"`, and `"phrase_type_t"` are all identifiers. However, `"8bf"` is not an identifier because it does not start with

T_IDENTIFIER referent	Case-sensitivity
Database name	insensitive
Object type name	insensitive
Enumeration name	insensitive
Enumeration constant name	sensitive

Table 3.1: Case-sensitivity of T_IDENTIFIERs

a letter or an underscore. Neither is `bf@foo.com` an identifier, because it does not consist solely of letters, underscores, and numbers.

Whether a T_IDENTIFIER is case-sensitive depends on what it stands for (i.e., what its “referent” is). See table 3.1 for a description.

4. A T_INTEGER is any sequence of one or more digits (0-9). For example, “0”, “42”, and “747” are all integers.
5. A T_STRING is one of two kinds:
 - (a) A T_STRING can start with a single quote (`'`), followed by zero or more characters which are not single quotes, and ending with another single quote (`'`). Such a string can contain newlines.
 - (b) A T_STRING can also start with a double quote (`"`), followed by zero or more characters, escape-sequences (see table 3.2), or newlines, and ending with a double quote (`"`).
6. White-space (spaces, newlines, and tabs) is ignored except in T_STRINGS.

Escape sequences	Meaning
<code>\n</code>	newline (ASCII 10)
<code>\t</code>	horizontal tab (ASCII 9)
<code>\v</code>	vertical tab (ASCII 11)
<code>\b</code>	backspace (ASCII 8)
<code>\a</code>	bell (ASCII 7)
<code>\r</code>	carriage-return (ASCII 13)
<code>\f</code>	form-feed (ASCII 12)
<code>\\</code>	slash (<code>\</code>) (ASCII 92)
<code>\?</code>	question-mark (<code>?</code>) (ASCII 63)
<code>\'</code>	single quote (<code>'</code>) (ASCII 39)
<code>\"</code>	double quote (<code>"</code>) (ASCII 34)
<code>\ooo</code>	Octal number (e.g., <code>\377</code> is 255)
<code>\xXX</code>	Hexadecimal number (e.g., <code>\xFF</code> is 255)

Table 3.2: Escape sequences in strings enclosed in double quotes.

3.1.4 Name-spaces

A name-space, in computer-terminology, is a language-domain with fixed borders within which names must be unique. *Within* a name-space, two different entities *cannot* be called by the same name without causing a name-clash. In other words, within a name-space, names must be unique. However, if two name-spaces are *orthogonal* to each other, then a name from one name-space *can* be the same as a name from the other name-space *without* causing a name-clash.

In MQL, the following name-spaces exist. They are all orthogonal to each other:

- Each object type forms a name-space with respect to its features. That is, a single object type cannot have two features with the same name, but different object types can have features with the same name. The two features with the same name need not even have the same feature type. This is because all name-spaces are orthogonal to each other.
- Each enumeration forms a name-space with respect to its constants. That is, a single enumeration cannot have two enumeration constants with the same name, but different enumerations can have enumeration constants with the same name. Since all name-spaces are orthogonal to each other, the two enumeration constants with the same name need not have the same integer value.
- Each database forms a global name-space with respect to object type names. That is, object type names must be globally unique within a database. However, since all name-spaces are orthogonal to each other, you can have features or enumeration constants which have the same name as an object type.

3.1.5 Top-level constraints on MQL syntax

The MQL engine can receive any number (greater than or equal to 1) of MQL statements. The only requirement is that each statement must end with the keyword "GO". This keyword acts as a delimiter between each statement. The last statement may also be terminated with "GO", but need not be. Single statements on their own need not be terminated with "GO" either.

If you connect to the MQL engine in daemon-mode, you must append the meta-level statement "QUIT" after the "GO" of the last statement.

3.2 Return types

3.2.1 Introduction

MQL is made up of statements, each of which either returns something or doesn't. If it returns something, there are two possibilities for what the return-type can be. It can be:

1. A table, or
2. A sheaf

The sheaf is explained in detail in section 4.4 on page 75. In section 3.2.3 below, we will treat the tables. But first, a word about the two output-formats available (plus the third option for getting data back).

3.2.2 Output-formats

The MQL engine gives you three options for using the results of an MQL query:

1. You can specify that you want XML output.
2. You can specify that you want output for displaying on a console.
3. You can use the datatype provided if your program is in the same process as the mql library.

If you use the mql(1) program for output, please see the manual page for how to choose the output kind.

3.2.3 Tables

The tables will look differently, depending on whether you choose XML-output or console-output. In the descriptions below, we will give abstract schemas for the tables, such as the following:

object_type_name : string	monad : monad_m	id_d : id_d
---------------------------	-----------------	-------------

This means that, in each row in the table, the first piece of data will be a string (called `object_type_name`), the second piece of data will be a `monad_m` (called `monad`), and the last piece of data will be an `id_d` (called `id_d`). And then the row stops. There will always be the same number of columns in each row.

A table of values may be empty, meaning it has no rows. In this case, there will still be a table heading with type-specifications.

Some MQL statements do not return a value. In this case, there will be no result, not even an empty table.

3.2.4 Atomic output-types in tables

The following types can get into a table and will be announced in the header of the table:

1. string
2. integer
3. boolean (true or false)
4. id_d

3.2.5 Other return values

A number of other values are also returned from each query:

1. A boolean indicating whether there were any compiler-errors.
2. A boolean indicating whether there were any database-errors.
3. An integer showing which stage of the compilation/interpretation we had come to when we exited the function (see table 3.3). In XML, this is a string as shown in the table, of the attribute “stage” attribute of the “error_stage” element.
4. A string carrying any error messages.

Stage	Value	XML string
None	0	none
Parsing	1	parse
Weeding	2	weed
Symbol-checking	3	symbol
Type-checking	4	type
Monads-checking	5	monads
Execution	6	exec

Table 3.3: Compiler stages. See `include/mql_execution_environment.h` for ready-made `#define` macros.

3.3 Database manipulation

3.3.1 CREATE DATABASE

3.3.1.1 Syntax

```
create_database_statement : "CREATE"  "DATABASE"  
    database_name opt_WITH_KEY  
;  
database_name : T_IDENTIFIER  
    | T_STRING  
;  
opt_WITH_KEY : /* Empty: No key is used. */  
    | "WITH"  "KEY"  T_STRING  
;
```

3.3.1.2 Example

```
CREATE DATABASE book_test  
GO
```

3.3.1.3 Explanation

The CREATE DATABASE statement creates and initializes a database. No text data is put into the database, and no object types are created, but the structures necessary for the EMdF engine to function are set in place. The user need not worry about these structures. Interested users are referred to [Relational-EMdF].

You must CREATE a database before you can USE it (see section 3.3.3). Alternatively, if you have a database that is already created but not initialized, you can use the INITIALIZE DATABASE statement (see Section 3.3.2).

If a transaction was in progress (see BEGIN TRANSACTION statement, section 3.4.1 on page 33), the transaction is automatically committed before the CREATE DATABASE statement is executed. Thus the user need not, cannot, and should not commit it or abort it.

The database name can be either a T_IDENTIFIER or a T_STRING. For MySQL and PostgreSQL, it must be a T_IDENTIFIER. For SQLite, it can be a T_STRING giving the filename (optionally including the full path) of the file in which the database is to be created. If no path is given, the file is created in the current working directory.

The optional "WITH KEY" syntax can be used on SQLite to send a key to SQLite's sqlite_open_encrypted API when opening the database. Note that this will not actually perform any encryption at all unless you obtain an encryption-enabled SQLite from somewhere, e.g., Dr. Hipp himself, the author of SQLite. See Section 2.8 on page 23 for more information.

3.3.1.4 Return type

There is no return value.

3.3.2 INITIALIZE DATABASE

3.3.2.1 Syntax

```
initialize_database_statement : "INITIALIZE"  "DATABASE"  
    database_name opt_WITH_KEY  
;  
database_name : T_IDENTIFIER  
    | T_STRING  
;  
opt_WITH_KEY : /* Empty: No key is used. */
```

```

        | "WITH" "KEY" T_STRING
;

```

3.3.2.2 Example

```

INITIALIZE DATABASE book_test
GO

```

3.3.2.3 Explanation

The INITIALIZE DATABASE statement initializes a database without creating it first. That is, the database must exist before issuing this statement. It simply creates all the meta-data necessary for having an Emdros database. This is useful on MySQL and PostgreSQL if you don't have privileges to create databases, but you do have privileges to create tables in an already-provided database. On SQLite, it is also useful, if you want to add Emdros information to an already-existing SQLite database. Other than not creating the database, this statement accomplishes the same things as the CREATE DATABASE statement (see Section 3.3.1 on the preceding page).

For the optional "WITH KEY" syntax, please see the CREATE DATABASE statement.

3.3.2.4 Return type

There is no return value.

3.3.3 USE DATABASE

3.3.3.1 Syntax

```

use_database_statement : "USE" [ "DATABASE" ]
                        database_name opt_WITH_KEY
;
database_name : T_IDENTIFIER
              | T_STRING
;
opt_WITH_KEY : /* Empty: No key is used. */
             | "WITH" "KEY" T_STRING
;

```

3.3.3.2 Example

```

USE book_test
GO

```

This is equivalent to

```

USE DATABASE book_test
GO

```

On SQLite:

```

USE DATABASE "c:\\EmdrosDBs\\mydb.db" /* On SQLite you can do this. */
GO

```

With a key:

```

/* On SQLite you can get encryption ****IF**** you have an
   encryption-enabled SQLite. */
USE DATABASE "c:\\Emdros\\MySecretDB.db"
/* The format and length of the key depends on your SQLite
   encryption implementation. This is just an example. */
WITH KEY "\\x45\\x98\\xbf\\x12\\xfa\\xc6"
GO

```

3.3.3.3 Explanation

Before you can start using a database you have CREATED (see section 3.3.1 on page 28) or INITIALIZED (see section 3.3.2 on page 28), you must connect to it using the USE DATABASE statement. The keyword “DATABASE” is optional and can be left out.

If a transaction was in progress (see BEGIN TRANSACTION statement, section 3.4.1 on page 33), the transaction is automatically committed before the USE DATABASE statement is executed. Thus the user need not, cannot, and should not commit it or abort it.

The database name can be either a T_IDENTIFIER or a T_STRING. For MySQL and PostgreSQL, it must be a T_IDENTIFIER. For SQLite, it can be a T_STRING giving the filename (optionally including the full path) of the file holding the database to be used. If no path is given, the file must be in the current working directory.

3.3.3.4 Return type

There is no return value.

3.3.4 DROP DATABASE

3.3.4.1 Syntax

```

drop_database_statement : "DROP" "DATABASE"
                        database_name
;
database_name : T_IDENTIFIER
              | T_STRING
;

```

3.3.4.2 Example

```

DROP DATABASE book_test
GO

```

3.3.4.3 Explanation

A previously CREATED database (see section 3.3.1) can be completely removed from the system using this statement. All data in the database is irretrievably lost, including all objects, all object types, and all enumerations.

If a transaction was in progress (see BEGIN TRANSACTION statement, section 3.4.1 on page 33), the transaction is automatically committed before the DROP DATABASE statement is executed. Thus the user need not, cannot, and should not commit it or abort it.

The database name can be either a T_IDENTIFIER or a T_STRING. For MySQL and PostgreSQL, it must be a T_IDENTIFIER. For SQLite, it can be a T_STRING giving the filename (optionally including the full path) of the file holding the database to be dropped. If no path is given, the file must be in the current working directory.

3.3.4.4 Return type

There is no return value.

3.3.5 VACUUM DATABASE

3.3.5.1 Syntax

```
vacuum_database_statement : "VACUUM" [ "DATABASE" ]
                             [ "ANALYZE" ]
;

```

3.3.5.2 Example

1. VACUUM DATABASE
GO
2. VACUUM DATABASE ANALYZE
GO

3.3.5.3 Explanation

On PostgreSQL, this statement vacuums the database using the “VACUUM” SQL statement. If the optional keyword “ANALYZE” is given, the statement issues a “VACUUM ANALYZE” statement. See the PostgreSQL documentation for what this does.

On MySQL, this statement issues OPTIMIZE TABLE queries for all object types. If the ANALYZE keyword is given, ANALYZE TABLE queries are issued as well.

On SQLite, this statement first deletes all redundant sequence info (compacting the sequence tables), then issues a VACUUM statement to SQLite.

The significance of this statement to Emdros development is that, when populating a database, things will speed up dramatically if the database is VACUUM’ed after every 1000 objects created, or so.

3.3.5.4 Return type

There is no return value.

3.3.6 DROP INDEXES

3.3.6.1 Syntax

```
drop_indexes_statement : "DROP" "INDEXES"
                        "ON" "OBJECT" ("TYPE" | "TYPES")
                        "[" object_type_to_drop_indexes_on "]"
;
object_type_to_drop_indexes_on : object_type_name | "ALL"
;
object_type_name : T_IDENTIFIER
;

```

3.3.6.2 Example

1. DROP INDEXES
ON OBJECT TYPES
[ALL]
GO

```

2. DROP INDEXES
   ON OBJECT TYPE
   [Word]
   GO

```

3.3.6.3 Explanation

Emdros creates indexes on the tables associated with object types when they are created. These indexes speed up retrieval, but slow down insertion. Therefore, if you are going to insert a large amount of objects, it is best to drop indexes on the object types you are going to modify (possibly all object types), then create the indexes again after you have inserted all objects.

This statement drops indexes that have previously been created. It has no effect if the indexes have been dropped already. If "ALL" is specified as the object type, then all object types in the current database will have their indexes dropped (if not dropped already).

The `manage_indices` program that comes with the Emdros distribution can be used to achieve the same effect.

Note that the choice between "TYPE" and "TYPES" is just syntactic sugar. It doesn't matter which you use.

If a feature has been declared WITH INDEX, this index is dropped. However, the feature will have its index recreated upon a CREATE INDEXES statement affecting that object type.

3.3.6.4 Return type

There is no return value.

3.3.7 CREATE INDEXES

3.3.7.1 Syntax

```

create_indexes_statement : "CREATE" "INDEXES"
                          "ON" "OBJECT" ("TYPE" | "TYPES")
                          "[" object_type_to_create_indexes_on "]"
;
object_type_to_create_indexes_on : object_type_name | "ALL"
;
object_type_name : T_IDENTIFIER
;

```

3.3.7.2 Example

```

1. CREATE INDEXES
   ON OBJECT TYPES
   [ALL]
   GO

2. CREATE INDEXES
   ON OBJECT TYPE
   [Word]
   GO

```

3.3.7.3 Explanation

Emdros creates indexes on the tables associated with object types when they are created. These indexes speed up retrieval, but slow down insertion. Therefore, if you are going to insert a large amount of objects,

it is best to drop indexes on the object types you are going to modify (possibly all object types), then create the indexes again after you have inserted all objects.

This statement creates indexes that have previously been dropped. It has no effect if the indexes are there already. If “ALL” is specified as the object type, then all object types in the current database will have their indexes created (if not there already).

The `manage_indices` program that comes with the Emdros distribution can be used to achieve the same effect.

Note that the choice between “TYPE” and “TYPES” is just syntactic sugar. It doesn’t matter which you use.

3.3.7.4 Return type

There is no return value.

3.4 Transactions

3.4.1 BEGIN TRANSACTION

3.4.1.1 Syntax

```
begin_transaction_statement : "BEGIN" "TRANSACTION"  
;
```

3.4.1.2 Example

```
BEGIN TRANSACTION  
GO
```

3.4.1.3 Explanation

On PostgreSQL, this statement begins a transaction if no transaction is in progress already. The return value is a boolean saying whether the transaction was started (true) or not (false). If this value is false, the user should not subsequently issue a `COMMIT TRANSACTION` or `ABORT TRANSACTION` statement. If this value is true, the user should issue either a `COMMIT TRANSACTION` or an `ABORT TRANSACTION` later.

On MySQL, this has no effect, and always returns false.

On SQLite, the behavior is the same as on PostgreSQL.

The transaction, if started, is automatically committed if a `CREATE DATABASE`, `USE DATABASE`, `DROP DATABASE` or `QUIT` statement is issued before a `COMMIT TRANSACTION` or `ABORT TRANSACTION` statement has been issued.

Also, the transaction is automatically committed if the connection to the database is lost, e.g., if the `mql(1)` program reaches the end of the MQL stream (e.g., an MQL script) and thus has to close down. Transactions are not maintained across invocations of the `mql(1)` program. The transaction is also committed if the `EMdFDB`, `CMQL_execution_environment`, or `CEmdrosEnv` object is destroyed.

3.4.1.4 Return type

A table with the following schema:

<code>transaction_started</code> : boolean
--

This table is empty if and only if there was a compiler error, i.e., if the syntax was not obeyed. The statement cannot fail with a database error. If no transaction was started, false is returned. If a transaction was started, true is returned.

3.4.2 COMMIT TRANSACTION

3.4.2.1 Syntax

```
commit_transaction_statement : "COMMIT" "TRANSACTION"  
;
```

3.4.2.2 Example

```
COMMIT TRANSACTION  
GO
```

3.4.2.3 Explanation

Commits the current transaction, if one is in progress. Has no effect if a transaction was not in progress. In such cases, false is returned.

If the commit failed, false is returned. If the commit succeeded, true is returned.

NOTE that this is slightly different from other statements which flag a DB error if unsuccessful. Here, no DB error is flagged, but false is returned in the table.

3.4.2.4 Return type

A table with the following schema:

transaction_committed : boolean

This table is empty if and only if there was a compiler error, i.e., if the syntax was not obeyed. The statement cannot fail with a database error. If no transaction was started when the COMMIT TRANSACTION statement was invoked, false is returned. If a transaction was started, and it was committed successfully, true is returned. If a transaction was started, but it was not committed successfully, false is returned.

3.4.3 ABORT TRANSACTION

3.4.3.1 Syntax

```
abort_transaction_statement : "ABORT" "TRANSACTION"  
;
```

3.4.3.2 Example

```
ABORT TRANSACTION  
GO
```

3.4.3.3 Explanation

Aborts the current transaction, if one is in progress. Has no effect if a transaction was not in progress. In such cases, false is returned.

If the abort failed, false is returned. If the abort succeeded, true is returned.

NOTE that this is slightly different from other statements which flag a DB error if unsuccessful. Here, no DB error is flagged, but false is returned in the table.

3.4.3.4 Return type

A table with the following schema:

```
transaction_aborted : boolean
```

This table is empty if and only if there was a compiler error, i.e., if the syntax was not obeyed. The statement cannot fail with a database error. If no transaction was started when the ABORT TRANSACTION statement was invoked, false is returned. If a transaction was started, and it was aborted successfully, true is returned. If a transaction was started, but it was not aborted successfully, false is returned.

3.5 Object type manipulation

3.5.1 CREATE OBJECT TYPE

3.5.1.1 Syntax

```
create_object_type_statement : "CREATE"
    [ "OBJECT" ] "TYPE"
    opt_range_type
    opt_monad_uniqueness_type
    "[" object_type_name
        [ feature_declaration_list ]
    "]"
;
opt_range_type:
    /* empty: Same as WITH MULTIPLE RANGE OBJECTS */
    | "WITH" "SINGLE" "MONAD" "OBJECTS"
    | "WITH" "SINGLE" "RANGE" "OBJECTS"
    | "WITH" "MULTIPLE" "RANGE" "OBJECTS"
;
opt_monad_uniqueness_type :
    /* empty: same as WITHOUT UNIQUE MONADS */
    | "HAVING" "UNIQUE" "FIRST" "MONADS"
    | "HAVING" "UNIQUE" "FIRST" "AND" "LAST" "MONADS"
    | "WITHOUT" "UNIQUE" "MONADS"
;
object_type_name : T_IDENTIFIER
;
feature_declaration_list : feature_declaration
    { feature_declaration }
;
feature_declaration : feature_name ":" feature_type
    [ default_specification ] ";"
    | feature_name ":" list_feature_type ";"
;
feature_type :
    "INTEGER" [with_index_specification]
    | "ID_D" [with_index_specification]
    | "STRING" [from_set_specification] [with_index_specification]
    | "ASCII" [from_set_specification] [with_index_specification]
    | T_IDENTIFIER /* For enumerations. */
;
list_feature_type :
    "LIST" "OF" "INTEGER"
```

```

        | "LIST" "OF" "ID_D"
        | "LIST" "OF" T_IDENTIFIER /* For enumerations */
;
with_index_specification :
    | "WITH" "INDEX"
    | "WITHOUT" "INDEX"
;
from_set_specification: "FROM" "SET"
;
default_specification : "DEFAULT" expression
;
expression : signed_integer /* integer and id_d */
    | T_STRING
    | T_IDENTIFIER /* enumeration constant */
;
signed_integer : T_INTEGER
    | "-" T_INTEGER
    | "NIL"
;

```

3.5.1.2 Examples

```

CREATE OBJECT WITH
WITH SINGLE MONAD OBJECTS
[Word
  surface: STRING FROM SET WITHOUT INDEX;
  lemma : STRING FROM SET WITH INDEX;
  psp : part_of_speech_t;
  parents : LIST OF id_d;
]
GO
CREATE OBJECT TYPE
[Clause
  parent : id_d;
  clause_type : clause_type_t default NC;
  functions : LIST OF clause_function_t; // An enumeration
  descendants : LIST OF ID_D;
]
GO

```

The latter creates an object type called “Clause” with four features: parent (Immediate Constituent of), whose type is `id_d`, and `clause_type`, which has the enumeration-type `clause_type_t` and the default value NC (which must be an enumeration constant in the `clause_type_t` enumeration). In addition, the two features “functions” and “descendants” are created, both of which are lists. “functions” is a list of enumeration constants drawn from the enumeration `clause_function_t`, whereas the “descendants” feature is a list of `id_ds`, which should then point to the descendants in the tree.

3.5.1.3 Explanation

This statement *creates an object type* in the meta-data repository of the current database. It starts out with the keywords “CREATE OBJECT TYPE”, followed by an optional clause which states whether the objects will be single-range or multiple-range (see below). After that comes a specification of the object type name and its features enclosed in square brackets. The `feature_declaration_list` is optional, so it is

possible for an object type to have no features.¹

Each `feature_declaration` consists of a feature name, followed by a colon, followed by a feature type, followed by an optional specification of the default value.

An `INTEGER`, `ID_D`, `STRING`, or `ASCII` feature can be declared “WITH INDEX”. This will put an index on the feature’s column. The default is not to add an index. This index will be dropped if a `DROP INDEXES` statement is issued for the object type (see Section 3.3.6 on page 31), but it will be recreated if a `CREATE INDEXES` statement is issued for the object type (see Section 3.3.7 on page 32). If a feature is an enumeration, it is usually not a good idea to create an index on the feature. This is because enumeration constants are usually few in number, and it is generally not a good idea to index columns that draw their values from a small pool of values, since this can lead to speed decreases ($O(N\log N)$) instead of $O(N)$). Therefore, the MQL language does not allow creating indexes on enumeration features. You can add them yourself, of course, if you like, with the backend’s corresponding SQL interface.

A `STRING` or `ASCII` feature can be declared “FROM SET”. The default is for it not to be from a set. This does *not* mean that the value of the feature *is* a set, but rather that the values are drawn FROM a set. Whenever an object is created or updated, and a feature is assigned which is declared “FROM SET”, the string value is first looked up in a special table that maps strings to unique integers. Then this unique integer is used in lieu of the string in the feature column. If the string does not exist in the separate table, it is added, with a unique integer to go with it, and that integer is used. If the string is there, then the integer already associated with the string is used. This gives a space savings (on MySQL and PostgreSQL), and often also a speed advantage, especially if used on all string features of a Word object type, and the words number many millions. SQLite may not see any speed or space savings advantage.

The specification of the default value (`default_specification`) consists of the keyword “DEFAULT”, followed by an expression. An expression is either a `signed_integer`, a string, or an identifier. The identifier must be an enumeration constant belonging to the enumeration which is also the type of the feature. The `signed_integer` is either a signed integer (positive or negative), or the keyword “NIL”, meaning the `id_d` that points to no object.

The feature “self” is implicitly created. It is an error if it is declared. The “self” feature is a computed feature which holds the unique object `id_d` of the object. See also 2.7.6 on page 21.

The difference between “ASCII” and “STRING” is that the user promises only to store 7-bit data in an ASCII string, whereas a `STRING` string may contain 8-bit data.²

In previous versions, you could specify a string length in parentheses after the `STRING` or `ASCII` keyword. As of version 1.2.0, all strings can have arbitrary length, with no maximum.³ The old syntax is still available, but is ignored.

If a feature is declared as a `LIST OF something`, that something has to be either `INTEGER`, `ID_D`, or an enumeration constant. Lists of strings are not supported. Also, you cannot declare a default value for a list – the default value is always the empty list.

An object type can be declared “WITH SINGLE MONAD OBJECTS”, “WITH SINGLE RANGE OBJECTS” or “WITH MULTIPLE RANGE OBJECTS”. The difference is that:

- An object type which has been declared “WITH SINGLE MONAD OBJECTS” can only hold objects which consist of a single monad (i.e., the first monad is the same as the last monad).
- An object type which has been declared “WITH SINGLE RANGE OBJECTS” can only hold objects which consist of a *single monad range*, i.e., there are no gaps in the monad set, but it consists of a single contiguous stretch of monads (possibly only 1 monad long).
- An object type which has been declared “WITH MULTIPLE RANGE OBJECTS” (the default) can hold objects which have arbitrary monad sets.

A single-monad object must consist of only 1 monad, e.g., {1}, {2}, {3}. A single-range object can consist of a single monad (e.g., { 1 }, { 2 }, { 3 }, etc.), or it can consist of a single interval (e.g., { 6-7 }, { 9-13 },

¹Strictly, this is not true, since all object types (except `pow_m`, `any_m`, and `all_m`) have at least one feature, namely the one called “self”. Please see section 2.7.6 on page 21 for more information.

²ASCII strings are stored exactly the same way as 8-bit STRINGS. This distinction is mostly obsolete.

³This is true for PostgreSQL (it is a TEXT). For MySQL, the maximum is 4294967295 ($2^{32} - 1$) characters (it is a LONGTEXT). On SQLite, it is a TEXT, but it is unknown how much this can hold.

{ 100-121 }, etc.). However, as soon as an object type needs to hold objects which can consist of more than one range (e.g., { 6-7, 9-13 }), then it must be declared WITH MULTIPLE RANGE OBJECTS. If neither is specified, then WITH MULTIPLE RANGE OBJECTS is assumed.

There is a speed advantage of using WITH SINGLE MONAD OBJECTS over WITH SINGLE RANGE OBJECTS, and again a speed advantage of using WITH SINGLE RANGE OBJECTS over WITH MULTIPLE RANGE OBJECTS. The latter is the slowest, but is also the most flexible in terms of monad sets.

In addition, and orthogonally to the range type, an object type can be declared “HAVING UNIQUE FIRST MONADS”, “HAVING UNIQUE FIRST AND LAST MONADS”, or “WITHOUT UNIQUE MONADS”. The difference is:

- An object type which has been declared “HAVING UNIQUE FIRST MONADS” can only hold objects which are unique in their first monad within the object type. That is, within this object type, no two objects may start at the same monad.
- An object type which has been declared “HAVING UNIQUE FIRST AND LAST MONADS” can only hold objects which are unique in their first monad and in their last monad (as a pair: You are allowed to have two objects with the same starting monad but different ending monads, or vice versa). That is, no two objects within this object type start at the same monad while also ending at the same monad. Note that for object types declared WITH SINGLE MONAD OBJECTS, a “HAVING UNIQUE FIRST AND LAST MONADS” restriction is upgraded to a “HAVING UNIQUE FIRST MONADS” restriction, since they are equivalent for this range type.
- An object type which has been declared “WITHOUT UNIQUE MONADS” (or which omits any of the “monad uniqueness constraints”) has no restrictions on the monads, other than those implied by the range type.

3.5.1.4 Return type

There is no return value.

3.5.2 UPDATE OBJECT TYPE

3.5.2.1 Syntax

```

update_object_type_statement : "UPDATE"
    [ "OBJECT" ] "TYPE"
    "[" object_type_name
        feature_update_list
    "]"
;
object_type_name : T_IDENTIFIER
;
feature_update_list : feature_update { feature_update }
;
feature_update : [ "ADD" ] feature_declaration
    | "REMOVE" feature_name ";"
;
feature_name : T_IDENTIFIER
;

```

3.5.2.2 References

All the foreign non-terminals are defined in section 3.5.1.

3.5.2.3 Example

```
UPDATE OBJECT TYPE
[Word
  ADD no_of_morphemes : integer;
  REMOVE surface_without_accents;
]
GO
```

This example ADDs the feature `no_of_morphemes` (being an integer), and REMOVEs the feature `surface_without_accents`.

3.5.2.4 Explanation

This statement *updates an object type*. It can either add a feature or remove an already-existing feature. When adding a new feature, the ADD keyword is optional. Other than that, it has exactly the same notation as for feature declarations under the CREATE OBJECT TYPE statement.

Removing a feature requires the REMOVE keyword, the feature name, and a semicolon.

Both additions and removals must be terminated with semicolon, even if the `feature_update` is the only `feature_update` in the list of `feature_updates`.

Note that the statement does not allow for *changing* the type of an already existing feature, only for adding or removing features.

3.5.2.5 Return type

There is no return value.

3.5.3 DROP OBJECT TYPE

3.5.3.1 Syntax

```
drop_object_type_statement : "DROP"
  [ "OBJECT" ] "TYPE"
  "[" object_type_name "]"
;
object_type_name : T_IDENTIFIER
;
```

3.5.3.2 Example

```
DROP OBJECT TYPE
[Sploinks]
GO
```

This example drops the object type “Sploinks” from the database.

3.5.3.3 Explanation

This statement drops an object type entirely from the database. This deletes not only the object type, but also all the objects of that object type, as well as the object type’s features. Enumerations which are used as a feature type are not dropped, however.

3.5.3.4 Return type

There is no return value.

3.6 Enumeration manipulation

3.6.1 CREATE ENUMERATION

3.6.1.1 Syntax

```
create_enumeration_statement : "CREATE"  
    ("ENUM" | "ENUMERATION")  
    enumeration_name "="  
    "{" ec_declaration_list "  
;  
enumeration_name : T_IDENTIFIER  
;  
ec_declaration_list : ec_declaration { "," ec_declaration }  
;  
ec_declaration : [ "DEFAULT" ]  
                ec_name  
                [ ec_initialization ]  
;  
ec_name : T_IDENTIFIER  
;  
ec_initialization : "=" signed_integer  
;
```

3.6.1.2 References

For a description of `signed_integer`, please see section 3.5.1 on page 35. Note, however, that `NIL` should not be used with enumerations.

3.6.1.3 Example

```
CREATE ENUMERATION  
phrase_type_t = { VP = 1, NP, AP,  
                 PP, default NotApplicable = -1 }  
GO
```

This particular statement creates an enumeration called “`phrase_type_t`” with the following constants and values:

Name	Value	Default
NotApplicable	-1	Yes
VP	1	No
NP	2	No
AP	3	No
PP	4	No

3.6.1.4 Explanation

This statement creates a new enumeration and populates it with enumeration constants.

If there is no declaration that has the “`default`” keyword, then the first one in the list becomes the default.

If the first declaration does not have an initialization, its value becomes 1. This is different from C and C++ enums, which get 0 as the first value by default.

If a declaration does not have an initialization, its values becomes that of the previous declaration, plus 1. This mimics C and C++ enums.

Label names must be unique within the enumeration. That is, you cannot have two constants with the same name in the same enumeration.

Values must also be unique within the enumeration. That is, you cannot have two different labels with the same value in the same enumeration. This is different from C/C++ enums, where two labels may have the same value.

3.6.1.5 Return type

There is no return value.

3.6.2 UPDATE ENUMERATION

3.6.2.1 Syntax

```

update_enumeration_statement : "UPDATE"
    ("ENUM" | "ENUMERATION")
    enumeration_name "="
    "{" ec_update_list  "}"
;
enumeration_name : T_IDENTIFIER
;
ec_update_list : ec_update { "," ec_update }
;
ec_update : [ "ADD" ] [ "DEFAULT" ]
    ec_name ec_initialization
    | "UPDATE" [ "DEFAULT" ] ec_name ec_initialization
    | "REMOVE" ec_name
;
ec_name : T_IDENTIFIER
;
ec_initialization : "=" signed_integer
;

```

3.6.2.2 References

For a description of `signed_integer`, please see section 3.5.1 on page 35. Note, however, that `NIL` should not be used with enumerations.

3.6.2.3 Example

```

UPDATE ENUMERATION
phrase_type_t = {
    ADD default Unknown = -99,
    REMOVE NotApplicable,
    UPDATE PP = 5,
    AdvP = 4
}
GO

```

This alters the table made in the example in section 3.6.1 to be like this:

Name	Value	Default
Unknown	-99	Yes
VP	1	No
NP	2	No
AP	3	No
AdvP	4	No
PP	5	No

3.6.2.4 Explanation

This statement updates the enumeration constants of an already existing enumeration. The user can specify whether to add, remove, or update an enumeration constant.

It is an error (and none of the updates will be executed) if the user REMOVES the default enumeration constant without specifying a new default.

Note that you are forced to specify values for all of the constants updated or added.

It is an error (and none of the updates will be executed) if the update would leave the enumeration in a state where two labels would have the same value. This is because an enumeration is effectively a one-to-one correspondence between a set of labels and a set of values.

It is the user's responsibility that the update leaves the database in a consistent state. For example, Emdros will not complain if you remove a constant with a given value without specifying a different constant with the same value, even if there are features that use this enumeration and have this value. This would mean that those feature-values could not be searched for, since there would be no label to look for. Neither would it be possible to get the feature values with GET FEATURES, since there would be no label to return.

3.6.2.5 Return type

There is no return value.

3.6.3 DROP ENUMERATION

3.6.3.1 Syntax

```
drop_enumeration_statement : "DROP"  
    ("ENUM" | "ENUMERATION")  
    enumeration_name  
    ;  
enumeration_name : T_IDENTIFIER  
    ;
```

3.6.3.2 Example

```
DROP ENUMERATION phrase_type_t  
GO
```

3.6.3.3 Explanation

This statement removes an enumeration altogether from the database, including all its enumeration constants.

It is an error (and impossible) to drop an enumeration which is in use by some object type.

3.6.3.4 Return type

There is no return value.

3.7 Segment manipulation

3.7.1 Introduction

Segments were present in Emdros up to and including version 1.1.12. After that, support for segments was removed.

A segment used to be an arbitrary, contiguous stretch of monads which was given a name. Objects could not be created which crossed the boundaries of a segment. You could restrict your search to within a single segment with `SELECT ALL OBJECTS`.

However, segments were found to be ugly baggage, cumbersome and not useful. Therefore, they were removed.

The `CREATE SEGMENT` statement is retained for backward compatibility.

3.7.2 CREATE SEGMENT

3.7.2.1 Syntax

```
create_segment_statement : "CREATE"    "SEGMENT"  
    segment_name  
    "RANGE"    "="    segment_range  
;  
segment_name : T_IDENTIFIER  
;  
segment_range : T_INTEGER "-" T_INTEGER  
;
```

3.7.2.2 Example

```
CREATE SEGMENT Old_Testament  
RANGE = 1 - 500000  
GO
```

This example used to create a segment named "Old_Testament" starting at monad 1 and ending at monad 500000.

Now it does nothing.

3.7.2.3 Explanation

This statement currently does nothing. It will fail with a database error.

3.7.2.4 Return type

There is no return value.

3.8 Querying

3.8.1 SELECT (FOCUS|ALL) OBJECTS

3.8.1.1 Syntax

```
select_objects_statement : select_clause  
    in_clause  
    returning_clause  
    where_clause  
;  
/*  
* select_clause  
*/  
select_clause : "SELECT"    focus_specification    [ "OBJECTS" ]  
;  
focus_specification : "FOCUS" | "ALL"
```

```

;
/*
 * in-clause
 */
in_clause : "IN"   in_specification
          | /* empty = all_m-1 */
;
in_specification : monad_set
                 | "ALL" /* = all_m-1 */
                 | monad_set
                 | T_IDENTIFIER /* Named arbitrary monad set */
;
monad_set : "{"   monad_set_element_list   "}"
;
monad_set_element_list : monad_set_element
                       { ", "   monad_set_element }
;
monad_set_element : T_INTEGER
                  | T_INTEGER   "-"   T_INTEGER
                  | T_INTEGER   "-"   /* From T_INTEGER to "practical infinity"
                                       (i.e., MAX_MONAD). */
;
/*
 * returning-clause
 */
returning_clause : /* Empty: Return full sheaf */
                 | "RETURNING" "FULL" "SHEAF"
                 | "RETURNING" "FLAT" "SHEAF"
                 | "RETURNING" "FLAT" "SHEAF" "ON"
                 object_type_name_list
;
object_type_name_list :
  object_type_name { ", " object_type_name }
;
/*
 * where-clause
 */
where_clause : "WHERE" mql_query
;

```

3.8.1.2 References

For the `mql_query` non-terminal, please see section 4.3 on page 72.

3.8.1.3 Example

```

SELECT ALL OBJECTS
IN { 1-4, 5, 7-9 }
RETURNING FULL SHEAF
WHERE
[Word lexeme = ">RY/"]
GO

```

3.8.1.4 Explanation

This statement is a front-end to the MQL Query-subset (see chapter 4 starting on page 68).

The parameters to an MQL query are:

1. A universe U,
2. A substrate Su, and
3. A topograph.

The universe U is a contiguous stretch of monads. The search is restricted only to include objects which are wholly contained within this universe (i.e., which are `part_of`⁴ the universe).

The substrate is used to further restrict the search. For there is the additional requirement that all objects found must be wholly contained within (i.e., `part_of`) the substrate as well. The substrate must be `part_of` the universe. Mathematically speaking, the substrate is the set intersection of whatever was in the IN clause and `all_m-1` (i.e., the set of all monads in the database).

The topograph is what is specified as `mql_query` in the above grammar.

The IN-specification tells the query-engine what the substrate Su should be. There are three choices:

1. Specify an explicit monad set like “{ 1-3000, 7000-10000 }”
2. Specify a named arbitrary monad set (see section 2.7.13 on page 23).
3. Leave it blank. This means that the substrate is calculated as `all_m-1` (i.e., all of the monads in the database; see [Standard-MdF] or [Doedens94] or page 20 in this Programmer’s Guide.)

The universe U is then calculated as all the monads between the first and last monads of the substrate.

The difference between the `RETURNING FULL SHEAF` and the `RETURNING FLAT SHEAF` clause is that the latter applies the “flatten” operator to the sheaf before returning it, whereas the former does not. If the `returning_clause` clause is empty, it means the same thing as `RETURNING FULL SHEAF`. If the `RETURNING FLAT SHEAF` has an `ON` appendix with a list of object type names, then the two-argument flatten operator is applied using this list of object type names. See section 4.4.7 on page 77 for an explanation of flat sheaves and the “flatten” operator.

3.8.1.5 Monad set

The explicit monad set in the IN clause, if given, must consist of a comma-separated list of monad-set-elements enclosed in curly braces. A monad-set-element is either a single integer (referring to a single monad) or a range consisting of two integers (referring to a range of monads). The monad-set-elements need not occur in any specific order, and are allowed to overlap. The result is calculated by adding all the monads together into one big set. The ranges of monads must, however, be monotonic, i.e., the second integer must be greater than or equal to the first.

3.8.1.6 Return type

A sheaf, either full or flat: All retrieved objects are included, but those objects that had the `focus` modifier in the query are flagged as such. Please see 4.4 on page 75 for an explanation of the sheaf. Appendix B on page 108 gives the grammar for the console-sheaf. Please see section 4.4.7 on page 77 for an explanation of the flat sheaf.

⁴For `part_of`, see section 2.7.7 on page 21.

3.8.2 SELECT OBJECTS AT

3.8.2.1 Syntax

```
select_objects_at_statement : "SELECT"    [ "OBJECTS" ]
    "AT"    single_monad_specification
    "["    object_type_name    "]"
;
single_monad_specification : "MONAD"    "="    T_INTEGER
;
object_type_to_find : object_type_name
;
object_type_name : T_IDENTIFIER
;
```

3.8.2.2 Example

```
SELECT OBJECTS
AT MONAD = 3406
[Clause]
GO
```

This example selects all those objects of type Clause which contain monad 3406.

3.8.2.3 Explanation

This statement returns a table containing the object `id_ds` of all the objects of the given type which have the monad specified, i.e., which contain the given monad.

The result is a table with one column, namely "id_d". Each row represents one object, where the `id_d` is its object `id_d`.

3.8.2.4 Return type

A table with the following schema:

id_d: id_d

On failure, this table is empty. Note, however, that the table can also be empty because there were no objects of the given type having the given monad. This is not an error.

3.8.3 SELECT OBJECTS HAVING MONADS IN

3.8.3.1 Syntax

```
select_objects_having_monads_in_statement :
    "SELECT" "OBJECTS"
    "HAVING" "MONADS" "IN"
    monad_set
    "["    object_type_to_find    "]"
;
object_type_to_find : object_type_name | "ALL"
;
object_type_name : T_IDENTIFIER
;
```

3.8.3.2 References

For the `monad_set` non-terminal, see section 3.8.1 on page 43.

3.8.3.3 Example

```
SELECT OBJECTS
HAVING MONADS IN { 23-45, 68, 70, 87-93 }
[Clause]
GO
```

```
SELECT OBJECTS
HAVING MONADS IN { 1, 5-7, 103-109 }
[ALL]
GO
```

```
SELECT OBJECTS
HAVING MONADS IN { 23 }
[Word]
GO
```

3.8.3.4 Explanation

This statement returns the object types and object `id_ds` of the objects that have at least one monad in the monad set specified. If “ALL” is specified as the object type, then this is done for all object types in the database. If a specific object type is specified, then that object type is used.

The returned table has one row for each object. Each object is represented only once. The monad in each row is guaranteed to be from the set of monads specified, and is guaranteed to be from the object in the row.

This statement is useful in much the same way that the `SELECT OBJECTS AT` statement is useful. It can be used, e.g., for getting `id_ds` of objects that must be displayed as part of the results of a query, but which are not in the query results. This statement can also be used like the `SELECT OBJECTS AT` statement by simply making the monad set a singleton set with only one monad. This statement has the advantage, however, that one can ask for all object types. This enables one to access objects which one knows might be there, but of which one does not know the object types. It also has the advantage of being much faster than a series of `SELECT OBJECTS AT` statements if one is looking for objects in more than one monad.

Note to programmers: If you want to get objects not from all object types but from only a subset of all object types, the easiest thing is to issue the required number of copies of the statement with `GO` in between, varying only the object type. That way, if you are using the `mql` program as a proxy for the `MQL` engine, you don't incur the overhead of starting and stopping the `mql` program.

3.8.3.5 Return type

<code>object_type_name</code> : string	<code>monad</code> : monad_m	<code>id_d</code> : id_d
--	------------------------------	--------------------------

On failure, this table is empty. Note, however, that the table can also be empty if the command were successful, if there were no objects that had at least one monad in the monad set specified.

3.8.4 GET OBJECTS HAVING MONADS IN

3.8.4.1 Syntax

```
get_objects_having_monads_in_statement :
  "GET" "OBJECTS"
  "HAVING" "MONADS" "IN"
  gohmi_monad_set
  "[ " object_type_name
      [gohmi_feature_retrieval]
  "]"
```

```

;
gohmi_monad_set : "ALL" /* all_m-1 */
                | monad_set
;
gohmi_feature_retrieval : "GET" feature_list
                        | "GET" "ALL"
;
feature_list : feature_name { "," feature_name }*
;
object_type_name : T_IDENTIFIER
;

```

3.8.4.2 References

For the monad_set non-terminal, see section 3.8.1 on page 43.

3.8.4.3 Example

```

GET OBJECTS
HAVING MONADS IN { 23-45, 68, 70, 87-93 }
[Clause]
GO

```

```

GET OBJECTS
HAVING MONADS IN { 1, 5-7, 103-109 }
[Phrase GET phrase_type, function]
GO

```

```

GET OBJECTS
HAVING MONADS IN { 23 }
[Word GET ALL]
GO

```

3.8.4.4 Explanation

This statement returns the objects of the given object type that have at least one monad in the monad set specified. The table returned has one row for each monad set element in the set of monads for each object returned.

This statement is useful in much the same way that the SELECT OBJECTS AT statement is useful. It can be used, e.g., for getting id_ds of objects that must be displayed as part of the results of a query, but which are not in the query results.

This is the preferred method for getting objects from the engine, rather than a sequence of SELECT OBJECTS HAVING MONADS IN, GET MONADS, and GET FEATURES. It is much faster than the combination of the three.

3.8.4.5 Return type

A flat sheaf is returned which contains the objects in question. See Section 4.4.7 on page 77 for more information.

3.8.5 SELECT OBJECT TYPES

3.8.5.1 Syntax

```

select_object_types_statement : "SELECT"

```

```

        [ "OBJECT" ] "TYPES"
;

```

3.8.5.2 Example

```

SELECT OBJECT TYPES
GO

```

3.8.5.3 Explanation

This statement returns a list of the names of all the object types available in the database.

3.8.5.4 Return type

A table with the following schema:

object_type_name: string

On failure, this table is empty.

3.8.6 SELECT FEATURES

3.8.6.1 Syntax

```

select_features_statement : "SELECT" "FEATURES"
    "FROM" [ [ "OBJECT" ] "TYPE" ]
    "[" object_type_name "]"
;
object_type_name : T_IDENTIFIER
;

```

3.8.6.2 Example

```

SELECT FEATURES
FROM OBJECT TYPE
[Phrase]
GO

```

3.8.6.3 Explanation

This statement returns a table with the features belonging to the given object type.

The type_name string in the result gives the type of the feature. It has the values as in table 3.4.

The default_value string in the result is a string representation of the default value. It must be interpreted according to the feature type.

The computed boolean in the result shows whether the feature is computed or not. Currently, the only computed feature is "self".

For lists, what is shown in the "default value" field is always "()" meaning "the empty list".

3.8.6.4 Return type

A table with the following schema:

feature_name: string	type_name: string	default_value: string	computed: boolean
----------------------	-------------------	-----------------------	-------------------

On failure, this table is empty. On success, the table cannot be empty, since every object type has the feature "self".

Type name	Meaning
integer	integer
id_d	id_d
list of integer	list of integer
list of id_d	list of id_d
list of <i>something else</i>	list of enumeration constants from the enumeration given
string	8-bit string of arbitrary length
ascii	7-bit (ASCII) string of arbitrary length
<i>everything else</i>	enumeration by the name given

Table 3.4: Possible type names in SELECT FEATURES

3.8.7 SELECT ENUMERATIONS

3.8.7.1 Syntax

```
select_enumerations_statement : "SELECT"
    "ENUMERATIONS"
    ;
```

3.8.7.2 Example

```
SELECT ENUMERATIONS
GO
```

3.8.7.3 Explanation

This statement returns a table with the names of all the enumerations available in the database.

3.8.7.4 Return type

A table with the following schema:

```
enumeration_name: string
```

On failure, this table is empty. Note, however, that it can also be empty because there are no enumerations in the database yet.

3.8.8 SELECT ENUMERATION CONSTANTS

3.8.8.1 Syntax

```
select_enumeration_constants_statement : "SELECT"
    ("ENUM" | "ENUMERATION") "CONSTANTS"
    "FROM" [ ("ENUM" | "ENUMERATION" ) ]
    enumeration_name
    ;
enumeration_name : T_IDENTIFIER
    ;
```

3.8.8.2 Example

```
SELECT ENUMERATION CONSTANTS
FROM ENUMERATION phrase_types
GO
```

3.8.8.3 Explanation

This statement returns a table with the enumeration constants in a given enumeration.

Note that the syntax is made so that the query need not be as verbose as in the example just given. There is quite a lot of syntactic sugar⁵ in this statement.

3.8.8.4 Return type

A table with the following schema:

enum_constant_name: string	value : integer	is_default : boolean
----------------------------	-----------------	----------------------

On failure, this table is empty.

3.8.9 SELECT OBJECT TYPES USING ENUMERATION

3.8.9.1 Syntax

```
select_object_types_which_use_enum_statement : "SELECT"  
  [ "OBJECT" ] "TYPES"  
  "USING"  
  ("ENUM" | "ENUMERATION") enumeration_name  
  ;  
enumeration_name : T_IDENTIFIER  
  ;
```

3.8.9.2 Example

```
SELECT OBJECT TYPES  
USING ENUMERATION phrase_types_t  
GO
```

3.8.9.3 Explanation

This statement returns a table with the names of the object types which use a given enumeration. The rows of the table are not ordered. An object type uses an enumeration if at least one of its features is of the enumeration type.

3.8.9.4 Return type

A table with the following schema:

object_type_name: string

On failure, this table is empty. Note, however, that it can also be empty because there are no object types using the enumeration.

3.8.10 SELECT MIN_M

3.8.10.1 Syntax

```
select_min_m_statement : "SELECT" "MIN_M"  
  ;
```

⁵“Syntactic sugar” is a term used by computer-scientists for niceties in the grammar of a language which help the user in some way, usually so that they do not have to type as much as would otherwise be required. Here, it simply means that some of the keywords are optional, or have shorthand forms.

3.8.10.2 Example

```
SELECT MIN_M  
GO
```

3.8.10.3 Explanation

Returns the minimum monad in use in the database. The table returned has only one data row, namely the minimum monad. See section 2.7.12 on page 22 for more information.

3.8.10.4 Return type

A table with the following schema:

min_m : monad_m

On failure, this table is empty.

3.8.11 SELECT MAX_M

3.8.11.1 Syntax

```
select_max_m_statement : "SELECT" "MAX_M"  
;
```

3.8.11.2 Example

```
SELECT MAX_M  
GO
```

3.8.11.3 Explanation

Returns the maximum monad in use in the database. The table returned has only one data row, namely the maximum monad. See section 2.7.12 on page 22 for more information.

3.8.11.4 Return type

A table with the following schema:

max_m : monad_m

On failure, this table is empty.

3.8.12 SELECT MONAD SETS

3.8.12.1 Syntax

```
select_monad_sets_statement : "SELECT" "MONAD" "SETS"  
;
```

3.8.12.2 Example

```
SELECT MONAD SETS GO
```

3.8.12.3 Explanation

This statement returns a table listing the names of the monad sets stored in the database. These are the monad sets referred to in section 2.7.13 on page 23.

The monad set names come in no particular order.

3.8.12.4 Return type

monad_set_name : string

3.8.13 GET MONAD SETS

3.8.13.1 Syntax

```
get_monad_sets_statement : "GET" "MONAD" ("SET" | "SETS")
                           monad_sets_specification
;
monad_sets_specification : "ALL"
                          | monad_set_list
;
monad_set_list : monad_set_name { "," monad_set_name }
;
monad_set_name : T_IDENTIFIER
;
;
```

3.8.13.2 Example

```
GET MONAD SET My_research_collection
GO
GET MONAD SETS Historical_books, Former_prophets
GO
GET MONAD SETS ALL
GO
```

3.8.13.3 Explanation

This statement returns a table listing the monads of each of the monad sets named in the query. These monad sets are the arbitrary monad sets described in section 2.7.13 on page 23.

It doesn't matter whether you say "SET" or "SETS". This is purely syntactic sugar.

If "ALL" is given as the monad_sets_specification, then all monad sets are listed, in no particular order.

In the output, each monad set is represented in the same way as described in section 3.8.14.4 on the following page. Each monad set is guaranteed to appear in the table in one contiguous stretch, that is, monad sets are not interleaved. Moreover, the monad set elements of each monad set is sorted on mse_first, in ascending order.

3.8.13.4 Return type

monad_set_name : string	mse_first : monad_m	mse_last : monad_m
-------------------------	---------------------	--------------------

3.8.14 GET MONADS

3.8.14.1 Syntax

```
get_monads_statement : "GET" "MONADS"
                      "FROM" ("OBJECT" | "OBJECTS")
                      "WITH" id_ds_specification
                      "[" object_type_name "]"
;
object_type_name : T_IDENTIFIER
;
;
```

3.8.14.2 References

For a description of `id_ds_specification`, please see section 3.9.2 on page 58.

3.8.14.3 Example

```
GET MONADS
FROM OBJECTS
WITH ID_DS = 10342, 10344, 10383
[Clause]
GO
```

3.8.14.4 Explanation

This statement returns, for each object in the list of `id_ds`, a representation of its set of monads. The set is represented by maximal stretches of monads. For example, if an object consists of the monads { 1, 2, 4, 5, 6, 9, 11, 12 }, and its `id_d` is 10342, then the following will be in the results of the above example:

object_id_d : id_d	mse_first : monad_m	mse_last : monad_m
10342	1	2
10342	4	6
10342	9	9
10342	11	12

The “mse” in “mse_first” and “mse_last” stands for “Monad Set Element.” A monad set element consists of a starting monad and an ending monad (always greater than or equal to the starting monad). It represents all of the monads between the two borders, including the borders. An mse’s last monad is always greater than or equal to its first monad.

The mses in the list are always maximal. That is, there is a gap of at least one monad in between each of the MSEs.

In mathematical terms, suppose we have an MSE A. Then for all other MSEs B for the same object, it is the case that either $A.last + 1 < B.first$ or $B.last < A.first - 1$

The MSEs will come in no particular order.

See [Monad Sets] for more information on monad sets and the way Emdros treats them.

It does not matter whether you write “OBJECT” or “OBJECTS”: The choice is merely syntactic sugar.

There is no limit on how many `id_ds` can be specified. The algorithm will not balk at even many thousand `id_ds`, but it will, of course, take more time to get the monads of more objects.

3.8.14.5 Return type

A table with the following schema:

object_id_d : id_d	mse_first : monad_m	mse_last : monad_m
--------------------	---------------------	--------------------

3.8.15 GET FEATURES

3.8.15.1 Syntax

```
get_features_statement : "GET"
    ("FEATURE" | "FEATURES")
    feature_list
    "FROM" ("OBJECT" | "OBJECTS")
    "WITH" id_ds_specification
    "[" object_type_name "]"
;
/*
* feature_list
```

```

*/
feature_list : feature_name { "," feature_name }
;
feature_name : T_IDENTIFIER
;
object_type_name : T_IDENTIFIER
;

```

3.8.15.2 References

For a description of `id_ds_specification`, please see section 3.9.2 on page 58.

3.8.15.3 Example

```

GET FEATURES surface, psp
FROM OBJECTS WITH ID_DS = 12513,12514
[Word]
GO

```

3.8.15.4 Explanation

This statement returns a table containing feature-values of certain objects in the database.

Note how this is different from the “SELECT FEATURES” command. The “SELECT FEATURES” command queries an *object type* for a list of its *features*. The “GET FEATURES” command queries *objects* for the *values* of some of their features.

3.8.15.5 Return type

The return type is a table with a schema containing one string for each feature in the list of features. The order of the columns is that in the list of features. The first column in the table contains the object `id_d` involved in the row. Thus for n features, the number of columns will be $n + 1$.

The return type of each feature is the same as the type of the feature. The exact representation depends on whether the output is console output or XML output. For XML, see the DTD. For console output, see the examples below. Enumeration constants are shown as the enumeration constant label, not the integer value.

For list features, the value is a space-surrounded, space-delimited list of values. Integers and `ID_Ds` are given as integers; enumeration constant values as their constant names (e.g., “first_person”).

object_id_d: id_d	surface: string	psp: enum(psp_t)	number_in_corpus: integer	parent: id_d	parents: list_of_id_d
-------------------	-----------------	------------------	---------------------------	--------------	-----------------------

The table contains the objects in no particular order.

3.8.16 GET SET FROM FEATURE

3.8.16.1 Syntax

```

get_set_from_feature_statement : "GET"      "SET"
    "FROM"      "FEATURE"
    feature_name
    "[" object_type_name "]"
;
feature_name : T_IDENTIFIER
;
object_type_name : T_IDENTIFIER
;

```

3.8.16.2 Example

```
GET SET
FROM FEATURE lexeme
[Word]
GO
```

3.8.16.3 Explanation

This statement returns a table containing the set of existing feature-values for a feature declared FROM SET. See the CREATE OBJECT TYPE and UPDATE OBJECT TYPE statements on page 35 and 38 respectively for the syntax of the FROM SET declaration.

Note how this is different from GET FEATURES: The “GET FEATURES” command queries *objects* for the *values* of some of their features. The GET SET FROM FEATURE queries the *set* of existing values for a given feature, regardless of which objects have these values for this feature.

3.8.16.4 Return type

The return type is a table with a schema containing one string for each value in the set.

value: string

The order of the strings in the table is undefined.

3.9 Object manipulation

3.9.1 CREATE OBJECT FROM MONADS

3.9.1.1 Syntax

```
create_object_from_monads_statement : "CREATE"    "OBJECT"
    "FROM"    monad_specification
    [ with_id_d_specification ]
    object_creation_specification
;
/*
 * monad-specification
 */
monad_specification : "MONADS"    "="    monad_set
;
/*
 * with-id_d-specification
 */
with_id_d_specification : "WITH"    "ID_D"
    "="    id_d_const
;
id_d_const : T_INTEGER
    | "NIL"
;
/*
 * object-creation-specification
 */
object_creation_specification : "["
    object_type_name
    [ list_of_feature_assignments ]
    "]"
```

```

;
object_type_name : T_IDENTIFIER
;
list_of_feature_assignments : feature_assignment
                             { feature_assignment }
;
feature_assignment : feature_name "!=" expression ";"
                   | feature_name "!=" list_expression ";"
;
feature_name : T_IDENTIFIER
;
expression : signed_integer /* integer and id_d */
            | T_STRING
            | T_IDENTIFIER /* enumeration constant */
;
list_expression : "(" [list_values] ")"
;
list_values : list_value { "," list_value }
;
list_value : signed_integer
            | T_IDENTIFIER /* enumeration constant */
;

```

3.9.1.2 References

For a description of `monad_set`, please see section 3.8.1 on page 43. For a description of `signed_integer`, please see section 3.5.1 on page 35. Note, however, that `NIL` should be used only with features whose type is `id_d`.

3.9.1.3 Example

```

CREATE OBJECT FROM MONADS = { 1-2, 4-7 }
[Clause
  clause_type := NC;
  parent := 10033;
  descendants := (104546, 104547, 104549);
]
GO

CREATE OBJECTS FROM MONADS = { 35-37 }
WITH ID_D = 104546
[Phrase
  phrase_type := NP;
  parents := (104212, 104215, 104219);
]
GO

```

3.9.1.4 Explanation

This statement creates a new object from a specified set of monads.

In creating an object, four items of information are necessary:

1. The new `id_d`,
2. The object type,

3. The set of monads,
4. Any features that need non-default values.

This statement creates an object of type “object_type_name” using the monads and features, and optional id_d, given. All features not specified will be given default values.

If you specify an id_d with the “WITH ID_D” specification, the system first checks whether that object id_d is already in use. If it is, the creation fails. If it is not, that id_d is used. If you do not specify an id_d, a unique id_d is auto-generated.

Note that when using WITH ID_D, it is not recommended to run several concurrent processes against the same database which issue CREATE OBJECT or CREATE OBJECTS statements. Doing so may cause the auto-generated object id_d sequence to become invalid. However, several concurrent processes may safely issue CREATE OBJECT(S) statements if none of them use WITH ID_D.

Note that objects of the special object types all_m, any_m, and pow_m cannot be created.

3.9.1.5 Return type

A table with the following schema:

object_id_d: id_d

On success, there is always only one row in the table, namely the row containing the object id_d of the newly created object.

On failure, the table is empty.

3.9.2 CREATE OBJECT FROM ID_DS

3.9.2.1 Syntax

```

create_object_from_id_ds_statement : "CREATE"      "OBJECT"
    "FROM"      id_ds_specification
    [ with_id_d_specification ]
    object_creation_specification
;
/*
 * id_ds-specification
 */
id_ds_specification : ("ID_D" | "ID_DS")
    "="      id_d_list
;
id_d_list : id_d  { "," id_d }
;
id_d : id_d_const
;
id_d_const : T_INTEGER
    | "NIL"
;

```

3.9.2.2 References

For the non-terminals with_id_d_specification and object_creation_specification, please see section 3.9.1 on page 56.

3.9.2.3 Example

```
CREATE OBJECT FROM ID_DS = 10028, 10029
[Clause
  clause_type := NC;
  parent := 10033;
]
GO
```

3.9.2.4 Explanation

This statement creates a new object with the monads contained in the objects specified by their `id_ds`.

The `id_ds` specified are used only to calculate the set of monads to be used. This is calculated as the union of the set of monads of the objects with the `id_ds` specified. These `id_ds` can point to objects of any type, and it need not be the same type for all `id_ds`.

Note that there is a syntactic sugar-choice of whether to say “ID_DS” or “ID_D”.

Note that objects of the special object types `all_m`, `any_m`, and `pow_m` cannot be created.

See the “CREATE OBJECT FROM MONADS” (section 3.9.1 on page 56) for further explanation and warnings. Especially about concurrent use of `WITH ID_D`.

3.9.2.5 Return type

The return type is the same as for `CREATE OBJECT FROM MONADS` (section 3.9.1).

3.9.3 CREATE OBJECTS WITH OBJECT TYPE

3.9.3.1 Syntax

```
create_objects_statement : "CREATE" "OBJECTS"
                          "WITH" "OBJECT" "TYPE"
                          "[" object_type_name "]"
                          object_creation_list
;

object_creation_list : object_creation { object_creation }
;

object_creation : "CREATE" "OBJECT"
                 "FROM" monad_specification
                 [ with_id_d_specification ]
                 "["
                 [ list_of_feature_assignments ]
                 "]"
;

monad_specification : monad_specification
;

with_id_d_specification : "WITH" ID_D "=" ID_DS
;

list_of_feature_assignments : list_of_feature_assignments
;

monad_specification : monad_specification
;

with_id_d_specification : "WITH" ID_D "=" ID_DS
;

list_of_feature_assignments : list_of_feature_assignments
```

3.9.3.2 References

For the non-terminals `monad_specification`, `with_id_d_specification`, and `list_of_feature_assignments` please see section 3.9.1 on page 56.

3.9.3.3 Example

```
CREATE OBJECTS
WITH OBJECT TYPE [Phrase]
CREATE OBJECT
```

```

FROM MONADS = { 1-2 }
[
  phrase_type := NP;
  function = Subj;
]
CREATE OBJECT
FROM MONADS = { 3-7 }
[
  // Use default values for phrase_type and function
  // (probably VP/Pred in this fictive example)
]
CREATE OBJECT
FROM MONADS = { 4-7 }
WITH ID_D = 1000000 // Assign specific ID_D
[
  phrase_type := NP;
  function := Objc;
]
GO

```

3.9.3.4 Explanation

This statement is for batch importing of objects. It is useful when populating databases, either from scratch or by adding large numbers of objects to an existing database. This statement is much faster than individual CREATE OBJECT statements.

The object type is specified only once, at the top. Note that no features can be assigned where the object type is specified: That comes later in the query, when each object is created.

Each object to be created must be given a monad set. The monad set follows the syntax specified in section 3.8.1 on page 43.

Optionally, an id_d can be specified. If an id_d is specified, it is the user's responsibility to ensure that the id_d assigned does not clash with another id_d in the database. This is mainly useful when dumping/restoring databases.

If no id_d is specified, a unique id_d is generated. This id_d is only guaranteed to be unique if no other objects are created with specific id_ds.

Note that when using WITH ID_D, it is not recommended to run several concurrent processes against the same database which issue CREATE OBJECT or CREATE OBJECTS statements. Doing so may cause the auto-generated object id_d sequence to become invalid. However, several concurrent processes may safely issue CREATE OBJECT(S) statements if none of them use WITH ID_D.

The feature-value assignments follow the same rules as for CREATE OBJECT FROM MONADS (see section 3.9.1 on page 56). If an object has a feature which is not assigned a value, the default value is used. The default value of a given feature can be specified when creating the object type, or when updating the object type (see section 3.5.1 on page 35 and section 3.5.2 on page 38).

A table is returned showing the number of objects created successfully. This number is valid even if the process failed half way through. In other words, if the process did not run to completion due to a DB error, the value in the return type will show how many objects, if any, were created successfully. This means that there is no way of knowing which object got which object id_d, a difference from the regular CREATE OBJECT statement.

3.9.3.5 Return type

A table with the following schema:

object_count: integer

On both success and failure, the table contains one row showing the number of objects created successfully.

3.9.4 UPDATE OBJECTS BY MONADS

3.9.4.1 Syntax

```
update_objects_by_monads_statement : "UPDATE"
    ("OBJECT" | "OBJECTS")
    "BY" monad_specification
    object_update_specification
;
/*
 * object-update-specification
 */
object_update_specification : "[" object_type_name
    list_of_feature_assignments
    "]"
;
object_type_name : T_IDENTIFIER
;
```

3.9.4.2 References

For the non-terminals `monad_specification` and `list_of_feature_assignments`, please see section 3.9.1 on page 56.

3.9.4.3 Example

```
UPDATE OBJECTS BY MONADS = { 1-2, 4-7, 8-20 }
[Clause
    clause_type := VC;
]
GO
```

3.9.4.4 Explanation

This statement finds all the objects of type `object_type_name` which are part_of the monads specified (i.e., they must be wholly contained within the monads specified), and updates their features according to the list of feature assignments.

Note that there is a syntactic sugar-choice of whether to say “OBJECTS” or “OBJECT”. This is because the user may know that only one object is to be found within the monads, in which case having to write “OBJECTS” would be intellectually irritating.

Note that objects of the special object types `all_m`, `any_m`, and `pow_m` cannot be updated.

The feature “self” cannot be updated.

3.9.4.5 Return type

A table with the following schema:

object_id_d: id_d

On success, the table contains one row for each updated object.

On failure, the table is empty.

3.9.5 UPDATE OBJECTS BY ID_DS

3.9.5.1 Syntax

```
update_objects_by_id_ds_statement : "UPDATE"  
    ("OBJECT" | "OBJECTS")  
    "BY"    id_ds_specification  
    object_update_specification  
    ;
```

3.9.5.2 References

For a description of `id_ds_specification`, see section 3.9.2 on page 58. For a description of `object_update_specification`, see section 3.9.4 on page 61.

3.9.5.3 Example

```
UPDATE OBJECTS BY ID_DS = 10028, 10029  
[Phrase  
    parent := 10034;  
]  
GO
```

3.9.5.4 Explanation

This statement updates all the objects of the given type with the given `id_ds`.

The `id_ds` should point to objects which are really of the given type. Otherwise, an error is issued.

Note that there is a syntactic sugar-choice between “OBJECTS” and “OBJECT.”

Note that objects of the special object types `all_m`, `any_m`, and `pow_m` cannot be updated.

The feature “self” cannot be updated.

3.9.5.5 Return type

The return type is the same as for UPDATE OBJECTS BY MONADS (section 3.9.4 on the previous page).

3.9.6 DELETE OBJECTS BY MONADS

3.9.6.1 Syntax

```
delete_objects_by_monads_statement : "DELETE"  
    ("OBJECT" | "OBJECTS")  
    "BY"    monad_specification  
    object_deletion_specification  
    ;  
/*  
 * object-deletion-specification  
 */  
object_deletion_specification : "["  
    object_type_name_to_delete  
    "]"  
    ;  
object_type_name_to_delete : object_type_name  
    | "ALL"  
    ;  
object_type_name : T_IDENTIFIER  
    ;
```

3.9.6.2 References

For a description of `monad_specification`, see section 3.9.1 on page 56.

3.9.6.3 Example

```
DELETE OBJECTS BY MONADS = { 1-20 }  
[Clause]  
GO
```

If “`object_name_to_delete`” is “ALL”, then all objects of all types which are at these monads are deleted:

```
DELETE OBJECTS BY MONADS = { 28901-52650 }  
[ALL]  
GO
```

3.9.6.4 Explanation

This command deletes all the objects of type `object_type_name` which are part_of the set of monads specified.

3.9.6.5 Return type

A table with the following schema:

object_id_d: id_d

On success, the table contains one row for each deleted object.
On failure, the table is empty.

3.9.7 DELETE OBJECTS BY ID_DS

3.9.7.1 Syntax

```
delete_objects_by_id_ds_statement : "DELETE"  
  ("OBJECT" | "OBJECTS")  
  "BY"   id_ds_specification  
  object_deletion_specification  
  ;
```

3.9.7.2 References

For a description of `id_ds_specification`, please see section 3.9.2 on page 58. For a description of `object_deletion_specification`, please see section 3.9.6 on page 62.

3.9.7.3 Example

```
DELETE OBJECTS BY ID_DS 10028, 10029  
[Phrase]  
GO
```

3.9.7.4 Explanation

This statement deletes objects by their `id_ds`. Note that you cannot write “ALL” for `object_deletion_specification`. The `id_ds` given should point to objects of the type given.

3.9.7.5 Return type

The return type is the same as for DELETE OBJECTS BY MONADS (section 3.9.6).

3.10 Monad manipulation

3.10.1 MONAD SET CALCULATION

3.10.1.1 Syntax

```
monad_set_calculation_statement : "MONAD" "SET"
    "CALCULATION"
    monad_set_chain
;
monad_set_chain : monad_set
    { monad_set_operator monad_set }
;
monad_set_operator : "UNION"
    | "DIFFERENCE"
    | "INTERSECT"
;
```

3.10.1.2 References

For a description of monad_set, please see section 3.8.1 on page 43.

3.10.1.3 Example

```
// Produces { 1-10 }
MONAD SET CALCULATION
{ 1-5, 7-8 }
UNION
{ 5-10 }
GO

// Produces { 2-5, 22-24 }
MONAD SET CALCULATION
{ 1-10, 20-30, 50-60 }
INTERSECT
{ 2-5, 22-24 }
GO

// Produces { 1-4, 8-10 }
MONAD SET CALCULATION
{ 1-10 }
DIFFERENCE
{ 5-7 }
GO

// Produces { 2-3, 5-6, 10-12 }
MONAD SET CALCULATION
{ 1-3, 5-9 }
INTERSECT
{ 2-6 }
UNION
```

```
{ 10-12 }  
GO
```

3.10.1.4 Explanation

This statement is for performing set-operations on sets of monads. The three standard set operations “union,” “intersect,” and “difference” are provided.

The return value is a representation of the resulting set of monads along the same lines as for the GET MONADS statement (see section 3.8.14).

The MSEs (see section 3.8.14) are listed in ascending order.

You can specify as many sets of monads as you want. The operations are done in succession from the first to the last set of monads. For example, in the last example above, the intersection is done first, and the union is done on the result of the intersection.

You can also specify only one set of monads, with no set operator. This is useful for creating a sorted, normalized set of monads from a number of different MSEs.

Note that this statement does not manipulate the stored arbitrary monad sets described in section 2.7.13 on page 23.

3.10.1.5 Return type

A table with the following schema:

mse_first : monad_m	mse_last : monad_m
---------------------	--------------------

3.10.2 CREATE MONAD SET

3.10.2.1 Syntax

```
create_monad_set_statement : "CREATE" "MONAD" "SET"  
    monad_set_name  
    "WITH" "MONADS" "=" monad_set  
;  
monad_set_name : T_IDENTIFIER  
;
```

3.10.2.2 References

For the monad_set non-terminal, please see section 3.8.1 on page 43.

3.10.2.3 Example

```
CREATE MONAD SET  
My_research_collection  
WITH MONADS = { 1-10394, 14524-29342, 309240-311925 }  
GO
```

3.10.2.4 Explanation

This statement creates an arbitrary monad set in the database. These monad sets are the ones described in section 2.7.13 on page 23.

3.10.2.5 Return type

There is no return value.

3.10.3 UPDATE MONAD SET

3.10.3.1 Syntax

```
update_monad_set_statement : "UPDATE" "MONAD" "SET"  
    monad_set_name  
    ("UNION" | "INTERSECT" | "DIFFERENCE" | "REPLACE")  
    (monad_set | monad_set_name)  
;  
monad_set_name : T_IDENTIFIER  
;
```

3.10.3.2 References

For the monad_set non-terminal, please see section 3.8.1 on page 43.

3.10.3.3 Examples

```
// Adds the specified monad set to "Historical_books"  
UPDATE MONAD SET  
Historical_books  
UNION  
{ 310320-329457 }  
GO  
// Remove the specified monad set from "Historical_books"  
UPDATE MONAD SET  
Historical_books  
DIFFERENCE  
{ 310320-329457 }  
GO  
// Intersects the monad set "My_research_collection"  
// with the monad set "My_experimental_collection"  
UPDATE MONAD SET  
My_research_collection  
INTERSECT  
My_experimental_collection  
GO  
// Replaces the monad set "Lamentations" with  
// the specified monad set  
UPDATE MONAD SET  
Lamentations  
REPLACE  
{ 380300-383840 }  
GO
```

3.10.3.4 Explanation

This statement is used to update an already-existing arbitrary monad set (see section 2.7.13 on page 23). Four operations are provided: set union, set intersection, set difference, and replacement. In all cases, the operation is done using two monad sets. The first set is the named set that is updated. The second set is either a set described in terms of monads, or the name of another arbitrary monad set.

The replacement operator effectively deletes the old set, replacing it with the new. Note, however, that this does not imply that the new is deleted – if you update one named monad set, replacing it with another named monad set, that other monad set is not deleted, but simply copied into the old monad set.

The other three operators are standard set-theoretic operators.

3.10.3.5 Return type

There is no return value.

3.10.4 DROP MONAD SET

3.10.4.1 Syntax

```
drop_monad_set_statement : "DROP" "MONAD" "SET"  
                           monad_set_name  
;  
monad_set_name : T_IDENTIFIER  
;
```

3.10.4.2 Example

```
DROP MONAD SET Historical_books  
GO
```

3.10.4.3 Explanation

This statement drops an arbitrary monad set (i.e., deletes it) from the database. These are the arbitrary monad sets described in section 2.7.13 on page 23.

3.10.4.4 Return type

There is no return value.

3.11 Meta-statements

3.11.1 QUIT

3.11.1.1 Syntax

```
quit_statement : "QUIT"  
;
```

3.11.1.2 Example

```
QUIT
```

3.11.1.3 Explanation

This causes the rest of the MQL stream not to be interpreted. It also causes the mql(1) program to quit after having executed this statement.

The QUIT statement can be used, e.g., if running the mql(1) program as a daemon through xinetd(8) or inetd(8), to end the connection.

The QUIT statement is special in that it does not need a "GO" keyword after it. You may supply the "GO" keyword if you wish, but it is not required.

If a transaction was in progress (see BEGIN TRANSACTION statement, section 3.4.1 on page 33), the transaction is automatically committed before the QUIT statement is executed.

3.11.1.4 Return type

There is no return value.

Chapter 4

MQL Query subset

4.1 Introduction

This chapter is an introduction to the query-subset of MQL for programmers. That is, it introduces the important subset of MQL in which you can express queries that find objects and gaps in interesting environments, with specified interrelations, and with specified feature-values.

An easier-to-read MQL Query Guide is available from the Emdros website, or with the Emdros source-code in the doc/ directory (see [MQLQueryGuide]).

First, we give an informal introduction to MQL by means of some examples (4.2). Then we give a complete overview of the syntax of the MQL query-subset (4.3). Then we explain the sheaf, which is the data-structure that an MQL query-query returns (4.4). Then we explain what a Universe and a Substrate are, since they are important in understanding how a query works (4.5). After that, we explain two important properties of mql queries, namely consecutiveness and embedding (4.6). After that, we give detailed explanations of the blocks of the MQL query-subset, which are the “building blocks” out of which a query is made (4.7). Finally, we explain how strings of blocks are written, and what they mean (4.8).

4.2 Informal introduction to MQL by means of some examples

4.2.1 Introduction

This section informally introduces the query-part of MQL by way of a number of examples. The example database which we will use is the same as in Doedens’ book, namely part of Melville’s “Moby Dick”:

“CALL me Ishmael. Some years ago - never mind how long precisely - having little or no money in my purse, and nothing particular to interest me on shore, I thought I would sail about a little and see the watery part of the world. It is a way I have of driving off the spleen, and regulating the circulation. Whenever I find myself growing grim about the mouth; whenever it is damp, drizzly November in my soul; whenever I find myself involuntarily pausing before coffin warehouses, and bringing up the rear of every funeral I meet; and especially whenever my hypos get such an upper hand of me, that it requires a strong moral principle to prevent me from deliberately stepping into the street, and methodically knocking people’s hats off - then, I account it high time to get to sea as soon as I can. [...]

“[...] By reason of these things, then, the whaling voyage was welcome; the great flood-gates of the wonder-world swung open, and in the wild conceits that swayed me to my purpose, two and two there floated into my inmost soul, endless processions of the whale, and, mid most of them all, one grand hooped phantom, like a snow hill in the air.”

Suppose that we have in this EMdF database the domain-dependent object types “paragraph”, “sentence”, and “word”, which correspond to paragraphs, sentences, and words of the text. And suppose that we add

to the object type “sentence” the feature “mood,” which draws its values from the enumeration type { imperative, declarative }. And suppose that we add to the object type “word” the features “surface” (which gives the surface text of the word) and “part_of_speech” (which gives the part of speech of the word). The codomain of the feature “part_of_speech” on the object type “word” draws its values from the enumeration type { adjective, adverb, conjunction, determiner, noun, numeral, particle, preposition, pronoun, verb }. This hypothetical database will give the background for most of the examples in our informal introduction to MQL.

In the following, when we refer to an “MQL query”, we will mean the query-subset of MQL. That is, we abstract away from the database-manipulation-part of MQL and concentrate on the query-queries. In addition, we will abstract away from the required “SELECT (FOCUS|ALL) OBJECTS” syntax that must precede an MQL query-query.

4.2.2 topograph

An MQL query is called a topograph. Consider the following topograph:

```
[ sentence ]
```

This topograph retrieves a list of all sentence objects in the database.

4.2.3 features

A query can specify which features an object must have for it to be retrieved. For example, consider the following topograph:

```
[ word
  surface = "Ishmael" or part_of_speech = verb;
]
```

This topograph retrieves a list of all words which either have the surface “Ishmael”, or whose part of speech is “verb.”

4.2.4 object_block, object_block_first

There are several types of blocks. They are meant to come in a string of blocks, where each block in the string must match some part of the database in order for the whole string to match. Two such blocks are the `object_block` and the `object_block_first`.

Object blocks are the heart and soul of MQL queries. They are used to match objects and objects nested in other objects. An object block (be it an `object_block` or an `object_block_first`) consists of the following parts:

1. The opening square bracket, ‘[’.
2. An identifier indicating the object type of the objects which we wish to match (e.g., “phrase”).
3. An optional “object reference declaration.” A reference to this object can be declared with the “as” keyword, like “[word as w...”. Subsequent blocks can then refer to features of this object as “w.featurename” (see section 4.7.7 on page 89).
4. An optional keyword which can be either of “noretrieve”, “retrieve” or “focus”. The default, when it is not specified, is “retrieve”. The keyword “noretrieve” says as much as “I do not wish to retrieve this object, even if matched”. It is useful for specifying the context of what we really wish to retrieve. The keyword “focus” specifies that this object is to be retrieved (it implies “retrieve”), and also that, when sifting the sheaf for focus objects, this object must go into the result (see section 4.7.4 on page 82).

5. An optional keyword, “first” or “last”, which says as much as “this object must be first/last in the universe against which we are matching (see section 4.7.5 on page 83).
6. An optional Boolean expression giving what features need to hold true for this object for it to be retrieved (see section 4.7.6 on page 84). This boolean expression must be prefixed by one of the words “feature” or “features”. It makes no difference which is used – it is merely syntactic sugar.
7. An optional inner blocks which matches objects inside the object (see section 4.8.3).
8. The closing square bracket, ‘]’.

Note that only the first object block in a string of blocks can have the “first” keyword, and only the last object_block in a string of blocks can have the “last” keyword.

Consider the following topograph:

```
[sentence
  mood = imperative;
  [word noretrieve first
    surface = "CALL";
  ]
  [word]
]
```

This topograph retrieves the set of sentences which are imperative, and whose first word is “CALL”. Within each sentence in that set, we retrieve the second word, but not the first. The only sentence in our example database which qualifies is the first sentence.

4.2.5 power

The power construct is used to indicate that we allow some distance in between two blocks. A power construct must always stand between two other blocks, and can thus never be first or last in a query. It comes in three varieties:

- A “plain vanilla” power construct, syntactically denoted by two dots, “. . .”, and
- A power construct with a single, upper limit. The limit specifies the maximum monads that can intervene between the two surrounding blocks. It is denoted as e.g., “. . . < 5”, or “. . . <= 5”.
- A power construct with a compound min/max limit. The limit specifies the minimum and maximum monads that can intervene. It is denoted as, e.g., “. . . BETWEEN 1 AND 5”.

Consider the following topograph:

```
[sentence
  [word
    part_of_speech = preposition]
  .. < 4
  [word
    part_of_speech = noun]
  ..
  [word last
    surface = "world"]
]
```

This topograph retrieves a list of sentences which have a word that has part of speech preposition, followed by a word which has part of speech noun, and which is within 4 monads of the preposition, followed by the last word of the sentence, which must be “world”. Within that sentence, retrieve all the three words. The only sentence which qualifies is the second.

4.2.6 opt_gap_block

An `opt_gap_block` is used to match an optional gap in the text. It consists of:

1. The opening square bracket, '['.
2. The keyword "gap?".
3. An optional "noretrieve," "retrieve" or "focus." The default is "noretrieve". (See section 4.7.4 on page 82)
4. An optional `blocks` (see section 4.8.3 on page 94).
5. The closing square bracket, ']'.

The `opt_gap_block` matches gaps in the *substrate* against which we are matching. Thus if we look at the example in figure 2.1 on page 19, we can construct the following topograph:

```
[clause
  [clause_atom
    [word
      surface = "door,"
    ]
  ]
  [gap? noretrieve]
  [clause_atom noretrieve]
]
```

This retrieves all clauses which happen to have inside them a `clause_atom` which contains the word "door," followed by a gap, followed by a `clause_atom`. The gap and the second `clause_atom` are not retrieved. This would retrieve clause-1. The gap need not be there.

The default is for the result of an `opt_gap_block` not to be retrieved. Thus one needs to explicitly write "retrieve" if one wishes to retrieve the gap.

4.2.7 gap_block

A `gap_block` is used to match a gap in the text. It consists of:

1. The opening square bracket, '['.
2. The keyword "gap".
3. An optional "noretrieve," "retrieve" or "focus." The default is "noretrieve". (See section 4.7.4 on page 82).
4. An optional `blocks`. (See section 4.8.3 on page 94).
5. The closing square bracket, ']'.

The `gap_block` is analogous to the `opt_gap_block` in all respects except that there *must* be a gap in order for the query to match.

4.2.8 object references

An object reference is a name given to a previously retrieved object with the “as *identifier*” declaration. An object reference can then be used in subsequent comparisons with features of other objects. This is done by selecting the desired feature from the object reference by using dot-notation, as in the example below:

```
[word as w
  part_of_speech = article;
]
[word
  (part_of_speech = noun
   or part_of_speech = adjective)
  and case = w.case
  and number = w.number
  and gender = w.gender;
]
```

Assuming that the `word` object type has features `part_of_speech`, `case`, `number`, and `gender`, this topograph retrieves all pairs of words which satisfy the following conditions:

- The first word has part of speech “article”,
- The second word has part of speech “noun” or “adjective”, and
- Both words have the same case, number, and gender.

This concludes our gentle, informal introduction to MQL.

4.3 Syntax of `mql_query`

4.3.1 Introduction

The `mql_query` non-terminal is the entry-point for the MQL query-subset. It is used in the `WHERE` clause of the `SELECT (FOCUS|ALL) OBJECTS` statement (section 3.8.1 on page 43). In this section, we give the full grammar of the MQL query-subset. It is important that you take some time to read through the grammar. Subsequent sections will build on the bird’s-eye view given in this section.

4.3.2 Syntax

```
mql_query : topograph
;
topograph : blocks
;
blocks : block_string
;
block_string : block_str
              | block_str "OR" block_string
;
block_str : block
           | block "!" block_str
           | block block_str
           | block power block_str
;
block : opt_gap_block
```

```

        | gap_block
        | object_block
        | ("NOTEXIST" | "NOTEXISTS" ) object_block
        | object_block "*" [monads_set]
    ;
opt_gap_block : "["   "GAP?"
                [ gap_retrieval ]
                [ blocks ]
                "]"
;
gap_retrieval : "NORETRIEVE"
                | "RETRIEVE"
                | "FOCUS"
;
gap_block : "["   "GAP"
            [ gap_retrieval ]
            [ blocks ]
            "]"
;
object_block : "["   object_type_name
               [ object_reference_declaration ]
               [ retrieval ]
               [ firstlast ]
               [ feature_constraints ]
               [ feature_retrieval ]
               [ blocks ]
               "]"
;
object_reference_declaration : "AS" object_reference
;
object_reference : T_IDENTIFIER
;
retrieval : "NORETRIEVE"
            | "RETRIEVE"
            | "FOCUS"
;
firstlast : "FIRST"
            | "LAST"
            | "FIRST" "AND" "LAST"
;
feature_constraints : ffeatures
;
ffeatures : fterm
            | ffeatures "OR" fterm
;
fterm : ffactor
        | ffactor "AND" fterm
;
ffactor : "NOT" ffactor
         | "(" ffeatures ")"
         | feature_comparison
;
feature_comparison :
    feature_name comparison_operator value

```

```

    | feature_name "IN"    enum_const_set
    | feature_name "IN"    "(" list_of_integer ")"
    | feature_name "IN"    object_reference_usage
;
comparison_operator : "="
                   | "<"
                   | ">"
                   | "<>" /* not equal */
                   | "<=" /* less than or equal */
                   | "<=" /* less than or equal */
                   | ">=" /* greater than or equal */
                   | ">=" /* greater than or equal */
                   | "~" /* regular expression */
                   | "!~" /* inverted regular expression */
                   | "HAS" /* lhs: list; rhs: atomic value.
                           signifies list membership. */
;
list_of_integer : T_INTEGER { ",", T_INTEGER }*
;
value : enum_const
      | signed_integer
      | T_STRING
      | object_reference_usage
;
enum_const : T_IDENTIFIER
;
object_reference_usage : object_reference
                      | "." feature_name
;
enum_const_set : "(" enum_const_list ")"
;
enum_const_list : enum_const { ",", enum_const_list }
;
power : ".." [ restrictor ]
;
restrictor : "<" limit
           | "<=" limit
           | "BETWEEN" limit "AND" limit
;
limit : T_INTEGER /* non-negative integer, may be 0. */
;
feature_retrieval : "GET" feature_list
                  | /* empty: Don't retrieve any features */
;

```

4.3.3 References

For the `signed_integer` non-terminal, please see section 3.5.1 on page 35. For `feature_name`, see 3.5.2 on page 38. For `feature_list`, see section 3.8.15 on page 54. For `monad_set`, see section 3.8.1 on page 43.

4.4 The sheaf

4.4.1 Introduction

The sheaf is the data structure that is returned from an MQL query-query. The structure of the sheaf closely reflects the structure of the query on which it is based. This section is meant as reading for implementors of Emdros-systems, not for end-users.

The sheaf has a specific structure, which we will look at next. After that, we will take a look at the meaning of the structures of the sheaf.

4.4.2 Structure of the sheaf

A sheaf consists of the following element types:

1. Sheaf
2. Straw
3. Matched_object

4.4.2.1 What is a sheaf?

A sheaf *is* a list of straws.

4.4.2.2 What is a straw?

A straw *is* a list of matched_objects.

4.4.2.3 What is a matched_object?

A matched_object *is* one of the following:

1. (object id_d, focus boolean, sheaf, object type, set of monads, list of feature-values)
2. (object id_m, focus boolean, sheaf)

That is, a matched_object is an object id (either id_d or id_m), coupled with a boolean indicating whether the block that gave rise to the matched_object had the “FOCUS” modifier, coupled with a sheaf. If the matched_object is of the first kind, then additionally, the object type and the object’s set of monads are also available, and there is a (possibly empty) list of feature-values.

4.4.3 MQL is topographic

There is a correspondence between the way an MQL query is structured and the structure of the resulting sheaf. In fact, the two are isomorphic to some extent. Doedens, in [Doedens94], called this property “topographicity.” Thus a `blocks` gives rise to a sheaf, a `block_str` gives rise to a straw, and a `block` gives rise to a matched_object. Inside a `block`, there is an optional inner `blocks`, which again gives rise to an inner sheaf. Hence a matched_object contains a sheaf. The origin of this sheaf is the optional inner `blocks` in the `block` which gave rise to the matched_object.

Note that this description applies to “full sheaves.” Flat sheaves are a different matter. See section 4.4.7 on page 77 for a description of flat sheaves.

4.4.4 Meaning of `matched_object`

A `matched_object` is the result of one of the following matches:

1. An `object_block` against an object in the database.
2. An `opt_gap_block` against a gap.
3. A `gap_block` against a gap.

A `matched_object`'s first component is either an `id_d` or an `id_m`. If the `matched_object` is the result of a match against an `object_block` or an `object_block_first`, then the `id` will be an `id_d`. If the `matched_object` is the result of a match against a `gap_block` or an `opt_gap_block`, the `id` is an `id_m`.

The second component is a boolean indicating whether the "FOCUS" keyword was present on the block.

The third component is a sheaf.

As we will see later, a sheaf is the result of matching against a `blocks`. It so happens that there is an optional `blocks` inside each of the four kinds of block (in the list above). The sheaf inside the `matched_object` is the result of a match against this `blocks`, if present. If the `blocks` is not present, then the sheaf is simply an empty sheaf.

For example, the following topograph:

```
[word FOCUS]
```

will contain one `matched_object` for each word-object within the substrate of the topograph. The sheaf of each of these `matched_objects` will be empty, and the FOCUS boolean will be "true" because we specified the FOCUS keyword.

4.4.5 Meaning of `straw`

A `straw` is the result of one complete match of a `block_str`. That is, a `straw` is a "string" of `matched_objects` corresponding to the blocks in the `block_str` which we should retrieve (which we can specify with the ("FOCUS"|"RETRIEVE"|"NORETRIEVE") keyword triad).

For example, consider the following topograph:

```
[word
  surface = "the";
]
[word
  part_of_speech = noun;
]
```

This will return a sheaf with as many `straws` as there are pairs of adjacent words where the first is the word "the" and the second is a noun. Each `straw` will contain two `matched_objects`, one for each word.

4.4.6 Meaning of the sheaf

A sheaf is the result of gathering all the matchings of a `blocks` non-terminal. There are four places in the MQL grammar where a `blocks` non-terminal shows up:

1. In the `topograph`,
2. In the `object_block`,
3. In the `opt_gap_block`, and
4. In the `gap_block`.

The first is the top-level non-terminal of the MQL query-query grammar. Thus the result of an MQL query-query is a sheaf.

Each of the last three is some kind of block. Inside each of these, there is an optional `blocks`. The result of matching this `blocks` is a sheaf.

But a sheaf is a list of straws. What does that mean?

It means that a sheaf contains as many matches of the `block_str` making up the `blocks` as are available within the substrate and universe that governed the matching of the `blocks`.

A straw constitutes one matching of the `block_str`. A sheaf, on the other hand, constitutes all the matchings.

4.4.7 Flat sheaf

Most of the above description has applied to “full sheaves.” We now describe flat sheaves.

A “flat sheaf,” like a “full sheaf,” consists of the datatypes “sheaf,” “straw,” and “matched_object.” The difference is that a “matched_object” in a flat sheaf cannot have an embedded sheaf. This makes a flat sheaf a non-recursive datastructure.

A flat sheaf arises from a full sheaf by means of the “flatten” operator.

If “FullSheaf” is a full sheaf, then “flatten(FullSheaf)” returns a flat sheaf that corresponds to the full sheaf.

A flat sheaf contains the same `matched_objects` as its originating full sheaf. However, they are structured such that each straw in the flat sheaf contains only `matched_objects` of one object type. Each object type that is represented in the full sheaf results in one straw in the flat sheaf.

Thus a straw in a flat sheaf does not correspond to the matching of a `block_string`. Instead, it is a list of all the `matched_objects` of one particular object type in the originating full sheaf. All of the `matched_objects` in the full sheaf are represented in the flat sheaf, regardless of whether they represent the same object in the database.

The “flatten” operator is only applied to the output of an MQL query if the “RETURNING FLAT SHEAF” clause is given (see section 3.8.1 on page 43). The programmer of an Emdros application can also apply it programmatically.

There is a variant of the flatten operator which also takes a list of object type names, in addition to the full sheaf. Then only those object types which are in the list are put into the flat sheaf. If `L` is a list of object type names, and `FullSheaf` is a full sheaf, then `flatten(FullSheaf, L)` returns a flat sheaf with straws for only those object types which are in `L`. If `L` is empty, then this is interpreted as meaning that all object types in `FullSheaf` must go into the flat sheaf. In the this light, the single-argument flatten operator may be seen as being a special case of the two-argument flatten operator, with `L` being empty. That is, `flatten(FullSheaf)` is the same as `flatten(FullSheaf, [])`.

4.5 Universe and substrate

4.5.1 Introduction

Two concepts which we shall need when explaining the blocks in MQL are “Universe” and “Substrate.” In this section, we define and explain them.

4.5.2 Universe and substrate

A Universe is a contiguous set of monads. It always starts at a particular monad a and ends at another monad b , where $a \leq b$. In more everyday language, a Universe is a stretch of monads that starts at one monad and ends at another monad later in the database. The ending monad may be the same as the starting monad.

A Substrate, on the other hand, is an arbitrary set of monads. It may have gaps (see section 2.7.8 on page 21). That is, while a Substrate always begins at a certain monad a and always ends at another monad b , where $a \leq b$, it need not contain all of the monads in between.

A Universe always has an accompanying Substrate, and a Substrate always has an accompanying Universe. Their starting- and ending-monads are the same. That is, the first monad of the Universe is always the same as the first monad of the accompanying Substrate. And the last monad of the Universe is always the same as the last monad of the Substrate. So a Universe is a Substrate with all the gaps (if any) filled in.

See section 3.8.1.4 on page 45 for an explanation of how the initial substrate and universe are calculated for the query.

With that definition out of the way, let us proceed to describing, exemplifying, and explaining *blocks*.

4.6 Consecutiveness and embedding

Two important notions in the MQL query-subset are embedding and consecutiveness. If two blocks (be they object blocks or gap blocks) are consecutive in a query, it means that they will only match two objects or gaps which are consecutive with respect to the substrate. Likewise, a string of blocks (i.e., a `blocks`) which is embedded inside of a block of some sort will only match within the confines of the monads of the surrounding block.

For example, the following topograph:

```
[Word psp=article]
[Word psp=noun]
```

will match two adjacent (or consecutive) words where the first is an article and the second is a noun. The consecutiveness is calculated with respect to the current substrate (see section 2.7.10 on page 21).

Likewise, the following topograph:

```
[Clause
  [Phrase phrase_type = NP]
  [Phrase phrase_type = VP]
]
```

will match only if the two (adjacent) phrases are found *within the confines of* the monads of the surrounding Clause. In fact the monads of the surrounding clause serve as the substrate when matching the inner blocks.

4.7 Blocks

4.7.1 Introduction

Blocks are the heart and soul of MQL query-queries. They specify which objects and which gaps in those objects should be matched and/or retrieved. With object blocks, you specify which objects should be matched. With gap blocks, you specify whether a gap should be looked for.

In this section, we treat the four kinds of blocks in MQL in some detail. First, we describe and explain the two kinds of object block (Object blocks, 4.7.2). Then we treat the two kinds of gap blocks (Gap blocks, 4.7.3). Then we describe how to specify whether to retrieve a block's contents (Retrieval, 4.7.4). After that we describe how to specify that an object block should be either first or last in its enclosing `blocks` (First and last, 4.7.5). Then we describe and explain how to specify constraints on features (Feature constraints, 4.7.6). Then we describe object references, which are a way of referring to other objects in a query (Object references, 4.7.7). After that we wrap up the syntactic non-terminals dealing with blocks by describing the `block` (Block, 4.7.8) Finally, we have some closing remarks (Closing remarks, ??).

4.7.2 Object blocks

4.7.2.1 Introduction

Object blocks specify which objects should be matched. Therefore, they are quite important in MQL. With object blocks, it is also possible to specify whether or not matched objects should be retrieved. You can also

specify constraints on the features of the objects which should be matched; You can specify whether you want objects matched against a certain object block to be first or last in the string of blocks we are looking for at the moment; And finally, you can label objects matched in a query with object reference labels, so that those objects can be referred to later in the query (i.e., further down in the MQL query, and thus further on in the string of monads). In this subsection, we deal with the object blocks themselves, deferring the treatment of feature-constraints, first/last-specifications, and object references to later subsections.

First, we describe the syntax of object blocks, then we give some examples, and finally we give some explanatory information.

4.7.2.2 Syntax

```

object_block : "["  object_type_name
               [ object_reference_declaration ]
               [ retrieval ]
               [ firstlast ]
               [ feature_constraints ]
               [ feature_retrieval ]
               [ blocks ]
               "]"
;
object_type_name : T_IDENTIFIER
;
retrieval : "NORETRIEVE"
           | "RETRIEVE"
           | "FOCUS"
;
firstlast : "FIRST"
            | "LAST"
            | "FIRST"  "AND"  "LAST"
;
last : "LAST"
;
feature_retrieval : "GET" feature_list
                   | /* empty: Don't retrieve any features */
;

```

4.7.2.3 References

For `object_reference_declaration`, see section 4.7.7 on page 89. For `feature_constraints`, see section 4.7.6 on page 84. For `blocks`, see section 4.8.3 on page 94. For `feature_list`, please see section 3.8.15 on page 54.

4.7.2.4 Examples

1. [Clause]
2. [Phrase noretrieve first
 phrase_type = NP
]
3. [Clause first and last]
4. [Word as w focus last
 psp = noun and number = pl]

```

    GET surface, lexeme
  ]

5. [Clause
  [Phrase first
    phrase_type = NP and phrase_function = Subj
  ]
  [Phrase
    phrase_type = VP
    [Word
      psp = V
    ]
    [Phrase
      phrase_type = NP and phrase_function = Obj
    ]
  ]
]

6. [Sentence
  NOTEXIST [Word surface = "saw"]
]

```

4.7.2.5 Explanation

Firstly, it will be noticed that the first item after the opening bracket must always be an object type name. This is in keeping with all other parts of MQL where object type names are used.

Secondly, it will be noticed that all of the other syntactic non-terminals in the definition of the object blocks are optional.

The object reference declaration comes after the object type name and will be dealt with below (4.7.7 on page 89).

The specification of the retrieval comes after the object reference declaration and will be dealt with in another section (4.7.4 on page 82).

The specification of first/last-constraints comes after the specification of retrieval, and will also be dealt with in another section (4.7.5 on page 83).

It is possible to specify constraints on the features of objects. This is done in the `feature_constraints` non-terminal, which comes after the first/last-constraints. These constraints will be dealt with in a section below (4.7.6 on page 84).

A list of features can be given in the `feature_retrieval` clause. Their values for a given object are placed on the list of features in the `matched_object` in the sheaf.

The inner `blocks` syntactic non-terminal allows the writer of MQL queries the possibility of matching objects nested inside objects. Example 5 above shows several examples of this. Example 5 finds those clauses which have inside them first a phrase which is a Subject NP, then followed by a Phrase which is a VP, the first word inside of which is a verb, followed by an Object NP. Thus we have an object block (“Clause”) with an inner `blocks` (“NP followed by VP”), where inside the VP we have another `blocks` (V followed by NP).

The inner `blocks`, if present, must match if the object block is to match. When entering the inner `blocks`, the Substrate for that `blocks` becomes the monads of the enclosing object. Let us call that object O. The Universe for the inner `blocks` becomes the set of monads between and including the borders of the enclosing object (see section 2.7.9 on page 21), i.e., the stretch of monads between (and including) O.first and O.last. This is the same as the substrate, except with any gaps filled in.

If you want any objects or gaps inside the object block to be retrieved, then the retrieval of the enclosing object block must be either `retrieve` or `focus`. Since the default retrieval for object blocks is to retrieve them, this condition is satisfied if you write nothing for the retrieval.

An object, if it is to match against a given object block, must be part_of the Substrate that accompanies the enclosing blocks. This basically means that all of the monads in the object must also be in the substrate. See section 2.7.7 on page 21 for details of the part_of relation.

In addition, the object, if it is to match against a given object block, must satisfy the feature constraints, if any. See section 4.7.6 on page 84 for details.

You can optionally place the keyword “NOTEXIST” before the object block. This will result in matching those cases where the object block does not occur, and will result in a failed match where the object block does occur. This is most useful if you have some context, i.e., a surrounding context (e.g., a sentence which does not contain such and such a word, see example 6 above). You are allowed to have blocks before and after a NOTEXIST block. Let us say that there is a block before the NOTEXIST block. Then the Substrate within which the NOTEXIST block will be matched is the Substrate of the context, minus the monads from the beginning of the Substrate to the end of the MatchedObject matching the previous block. The Universe of the NOTEXIST block will be defined analogously on the Universe of the context.

The NOTEXIST block will have “zero width” with respect to consecutiveness: If it matches anything, the entire block_string fails. If it does not match, it is as though the NOTEXIST block had not been there, and any block after the NOTEXIST block will be attempted matched starting at the previous block’s last monad plus 1.

The NOTEXIST keyword acts as an “upwards export barrier” of object reference declarations. That is, you cannot “see” an object reference declaration outside of the NOTEXIST, only inside of it.

4.7.3 Gap blocks

4.7.3.1 Introduction

Gap blocks are used to match gaps in the substrate we are currently matching against. There are two kinds of blocks: plain gap blocks and optional gap blocks.

We start by defining the syntax related to gap blocks. We then give some examples of gap blocks. And finally, we provide some explanation.

4.7.3.2 Syntax

```
gap_block : "[" "GAP"
           [ gap_retrieval ]
           [ blocks ]
           "]"
;
opt_gap_block : "[" "GAP?"
                 [ gap_retrieval ]
                 [ blocks ]
                 "]"
;
gap_retrieval : "NORETRIEVE"
               | "RETRIEVE"
               | "FOCUS"
;

```

4.7.3.3 Examples

1. [gap]
2. [gap?]
3. [gap noretrieve]
4. [gap? focus]

```

5. [gap retrieve
    [Word retrieve
      psp = particle
    ]
  ]

```

4.7.3.4 Explanation

There are two differences between the two types of gap block: One is that the `gap_block` *must* match a gap in the substrate for the whole query to match, while the `opt_gap_block` *may* (but need not) match a gap in the substrate. The other is that the default retrieval of an `opt_gap_block` is `NORETRIEVE`, whereas the default retrieval of a `gap_block` is `RETRIEVE`. Otherwise, they are identical in semantics.

The retrieval will be dealt with more fully in the next section.

The inner `blocks`, if present, must match if the gap block is to match. When trying to match the inner `blocks`, both the `Universe` and the `Substrate` are set to the monads of the gap. So if the gap matches the monad-stretch `[a..b]`, then both the `Universe` and the `Substrate` for the inner `blocks` will be this stretch of monads.

The last point is important in example 5. Here the `Word` which we are looking for inside the gap will be looked for within the monads which made up the gap.

If you want any objects or gaps to be retrieved inside the gap (as in example 5 above, where we want to retrieve the `Word`), then the retrieval of the gap block must be either “retrieve” or “focus”.

4.7.4 Retrieval

4.7.4.1 Introduction

Retrieval is used in four places in the MQL grammar. Once for each of the two object blocks and once for each of the two gap blocks. In this section we describe the three kinds of retrieval, specify the default behavior, and provide a bit of explanation.

4.7.4.2 Syntax

```

retrieval : "NORETRIEVE"
           | "RETRIEVE"
           | "FOCUS"
;
gap_retrieval : "NORETRIEVE"
               | "RETRIEVE"
               | "FOCUS"
;

```

4.7.4.3 Examples

1. [Word focus
 psp = verb
]
2. [gap? retrieve]
3. [Phrase noretrieve]
4. [gap focus]
5. [Phrase retrieve]

```

    [Word focus
      psp = article
    ]
    [gap retrieve
      [Word focus
        psp=conjunction
      ]
    ]
  [Word focus
    psp = noun
  ]
]

```

4.7.4.4 Explanation

Retrieval has to do with two domains pertaining to objects and gaps:

1. Whether to retrieve the objects or gaps, and
2. Whether those objects or gaps should be in *focus*.

Whether to retrieve is straightforward to understand. If we don't retrieve, then the object or gap doesn't get into the sheaf. The sheaf is the data-structure returned by an MQL query. The object or gap (if the gap is not optional) must still match for the overall match to be successful, but the object or gap won't get into the sheaf if we don't retrieve.

When an object is in focus, that means your application has the opportunity to filter this object out specifically from among all the objects retrieved. Exactly how this feature is used (or not used) will depend on your application. When is this useful?

Recall that, for objects in an inner `blocks` to be retrieved (in an object block or a gap block), the enclosing object or gap must also be retrieved. Thus you might end up with objects in the sheaf which you don't really care about. The focus-modifier is a way of signaling special interest in certain objects or gaps. Thus you can specify exactly which objects should be of special interest to the application. In example 5 above,¹ the outer `Phrase` must be retrieved, because we wish to retrieve the inner objects and gaps. The inner `gap` must also be retrieved because we wish to retrieve the inner `Word`. The three `Words` are what we are really interested in, however, so we mark their retrieval as "focus".

If we specify "focus" as the retrieval, then that implies "retrieve". Thus we can't not retrieve an object which is in "focus". This makes sense. If you have registered a special interest in an object, that means you want to retrieve it as well.

The default for object blocks of both kinds, when no retrieval is specified, is to assume "retrieve". The default for gap blocks of both kinds, on the other hand, is "noretrieve".

4.7.5 First and last

4.7.5.1 Introduction

The object blocks have the option of specifying whether they should be first and/or last in their enclosing blocks.

4.7.5.2 Syntax

```

firstlast : "FIRST"
           | "LAST"
           | "FIRST" "AND" "LAST"
;

```

¹This construction actually does occur in at least one language, namely ancient Greek. It is due to post-positive particles and conjunctions such as "de", "gar", "men", and the like.

4.7.5.3 Examples

1. [Clause first and last]
2. [Phrase first]
3. [Clause
 [Phrase first]
 [Word last
 psp = verb
]
]

4.7.5.4 Explanation

In example 1, the clause must be both first and last in its surrounding `blocks`. In the second example, the phrase must merely be the first. In the third example, the Phrase must be first in the clause, followed by a word, which must be a verb, and which must be last. This can be realized, e.g., in verb-final languages.

What does it mean to be “first” and “last” in the enclosing `blocks`?

Again we must appeal to the notion of Universe and Substrate. Each `blocks` carries with it a Universe and a Substrate. Let us say that an object block must be first, and let us say that we are trying to match an object `O` against this object block. Let us call the substrate of the enclosing blocks “`Su`”. Then, for the object `O` to be first in the blocks means that `O.first = Su.first`. That is, the first monad of the object must be the same as the first monad of the Substrate.

Conversely, for an object `O` to be last in a `blocks`, means that `O.last = Su.last`. That is, the last monad of the object must be the same as the last monad of the Substrate.

4.7.6 Feature constraints

4.7.6.1 Introduction

Object blocks can optionally have feature constraints. The feature constraints are boolean (i.e., logical) expressions whose basic boolean building-blocks are “and”, “or”, and “not”. The things that are related logically are comparisons of features and values, i.e., a feature followed by a comparison-symbol (e.g., “=”), followed by a value. Parentheses are allowed to make groupings explicit.

In the following, we first define the syntax of feature constraints. We then make refer to other parts of this manual for details of certain non-terminals. We then give some examples, followed by explanations of those examples. We then give some explanation and elucidation on feature-constraints. We then describe the constraints on type-compatibility between the feature and the value. Finally we elaborate on comparison-operators.

4.7.6.2 Syntax

```
feature_constraints : ffeatures
;
ffeatures : fterm
           | ffeatures "OR" fterm
;
fterm : ffactor
       | ffactor "AND" fterm
;
ffactor : "NOT" ffactor
         | "(" ffeatures ")"
         | feature_comparison
;
;
```

```

feature_comparison :
    feature_name  comparison_operator  value
    | feature_name "IN"    enum_const_set
    | feature_name "IN"    "(" list_of_integer ")"
    | feature_name "IN"    object_reference_usage
;
comparison_operator : "="
                    | "<"
                    | ">"
                    | "<>" /* not equal */
                    | "<=" /* less than or equal */
                    | ">=" /* greater than or equal */
                    | "~" /* regular expression */
                    | "!~" /* inverted regular expression */
                    | "HAS" /* lhs: list; rhs: atomic value.
                             signifies list membership. */
;
list_of_integer : T_INTEGER { ",", T_INTEGER }*
;
value : enum_const
      | signed_integer
      | T_STRING
      | object_reference_usage
;
enum_const : T_IDENTIFIER
;
object_reference_usage : object_reference
                       "." feature_name
;
enum_const_set : "(" enum_const_list ")"
;
enum_const_list : enum_const { ",", enum_const_list }
;

```

4.7.6.3 References

For `signed_integer`, see section 3.5.1 on page 35. For object references, see the next section. For `feature_name`, see 3.5.2 on page 38.

4.7.6.4 Examples

1. [Word psp = noun]
2. [Word gender = neut or gender = fem]
3. [Word psp = adjective and not case = nominative]
4. [Phrase (phrase_type = NP
 and phrase_determination = indetermined)
 or phrase_type = AP
]
5. [Word as w
 psp = article

- ```

]
[Word
 psp = noun
 and case = w.case
 and gender = w.gender
 and number = w.number
]

6. [Word
 surface > "Aa" and surface ~ "[A-D]orkin"
]

7. [Word psp IN (verb, participle, infinitive)]

8. [Word psp = verb OR psp = participle OR psp = infinitive]

```

#### 4.7.6.5 Explanation of Examples

Example 1 above is the simple case where a feature (“*psp*”) is being tested for equality with a value (“*noun*”). Example 2 is more of the same, except the gender can either be neuter or feminine, and the feature constraint would match in both cases. Example 3 finds those words which are adjectives *and* whose case is *not* nominative. Example 4 finds either adjectival phrases or NPs which are indetermined.

Example 5 is an example of usage of object references. The first Word is given the “label” (or “object reference”) “*w*”. Then the second Word’s feature-constraints refer to the values of the features of the first Word, in this case making sure that case, number, and gender are the same.

Example 6 is an example of two different comparison-operators, “greater-than” and “regular-expression-match”.

Example 7 shows the comparison IN. It takes a comma-separated list of enumeration constant names in parentheses as its right-hand-side. The effect is the same as an OR-separated list of “=” feature-comparisons. So 7. and 8. are equivalent.

#### 4.7.6.6 Explanation

While the syntax may look daunting to the uninitiated, the system is quite straightforward. At the bottom, we have feature comparisons. These consist of a feature, followed by a comparison-operator (such as “=”), followed by a value. These feature-comparisons can be joined by the three standard boolean operators “and”, “or”, and “not”.

The precedence of the operators follows standard practice, i.e., “not” has highest precedence, followed by “and”, followed by “or”. Parentheses are allowed to make groupings explicit. That is, “and” “binds” more closely than “or” so that the interpretation of this expression:

$$f_1 = val_1 \text{ “and” } f_2 = val_2 \text{ “or” } f_3 = val_3$$

is the following:

$$(f_1 = val_1 \text{ “and” } f_2 = val_2) \text{ “or” } f_3 = val_3$$

Note that if you use “not” on a feature comparison, and if you have another feature comparison before it, then you must explicitly state whether the relationship between the two is “and” or “or”. Thus the following is illegal:

$$f_1 = val_1 \text{ “not” } f_2 = val_2$$

The following, however, would be legal:

$$f_1 = \text{val}_1 \text{ “and” “not” } f_2 = \text{val}_2$$

The “in” comparison-operator can only be used with a comma-separated list of enumeration constant names on the right-hand-side. The effect is the same as if all of the enumeration constants had been compared “=” to the feature, with “OR” between them.

#### 4.7.6.7 Type-compatibility

The feature and the value with which we compare both have a *type*. The type is, one of “integer”, “7-bit (ASCII) string”, “8-bit string”, “enumeration constant”, “id\_d”. Thus a type tells us how to interpret a value.

The types of the two values being compared must be *compatible*. Table 4.1 summarizes the type-compatibility-constraints.

| If value’s type is...       | Then feature’s type must be...                                                         |
|-----------------------------|----------------------------------------------------------------------------------------|
| enumeration constant        | The same enumeration as the value                                                      |
| (enumeration constant-list) | The same enumeration as all the values                                                 |
| signed_integer              | integer or id_d                                                                        |
| 7-bit or 8-bit string       | 7-bit or 8-bit string                                                                  |
| object reference usage      | The same type as the feature in the object reference usage, or a list of the same type |

Table 4.1: Type-compatibility-constraints

The 8-bit strings need not be of the same encoding.

#### 4.7.6.8 Comparison-operators

Table 4.2 summarizes the comparison-operators.

| op. | meaning                                                                             |
|-----|-------------------------------------------------------------------------------------|
| =   | Equality                                                                            |
| <   | Less-than                                                                           |
| >   | Greater-than                                                                        |
| <>  | Inequality (different from)                                                         |
| <=  | Less-than-or-equal-to                                                               |
| >=  | Greater-than-or-equal-to                                                            |
| ~   | Regular expression-match                                                            |
| !~  | Negated regular-expression-match                                                    |
| IN  | Member of a list of enum constants                                                  |
| HAS | List on left-hand-side, atomic value on right-hand-side. Signifies list membership. |

Table 4.2: Comparison-operators

**4.7.6.8.1 Inequality** The inequality-operator “<>” is logically equivalent to “not ... = ...”. The negated regular-expression-match “!~” is logically equivalent to “not ... ~ ...”.

**4.7.6.8.2 Equality** Equality is defined as follows: If the type is `id_d`, then both must be the same `id_d`. If the type is integer, then both must be the same number. If the type is string, then both must be byte-for-byte identical, and of the same length. If the type is enumeration, then both must have the same numerical value. That is, the enumeration constants must be the same, since an enumeration is a one-to-one correspondence between a set of labels and a set of values. If the type is a list, the two lists must be identical, i.e., consist of the same sequence of values.

**4.7.6.8.3 Less-than/greater-than** The four less-than/greater-than-operators use 8-bit scalar values for the comparison of strings. That is, it is the numerical value of the bytes in the strings that determine the comparison. In particular, the locale is not taken into consideration. For comparison of `id_ds`, the `id_ds` are treated as ordinary numbers, with `nil` being lower than everything else. For comparison of integers, the usual rules apply. For comparison of enumeration constants, it is the values of the enumeration constants that are compared, as integers.

**4.7.6.8.4 Regular expressions** There are two regular expression-comparison-operators (“`~`” and “`!~`”). They operate on 8-bit strings. That is, both the feature-type and the value against which the match is made must be 8-bit strings. The negated match matches everything that does not match the regular expression provided as the value.

The value that they are matched against must be a string.

The regular expressions are the same as in Perl 5. See section 1.4.2 on page 15 for details of where regular expression-support comes from. See <http://www.perl.com/> for details of Perl regular expressions.

Before version 1.2.0.pre46, regular expressions were anchored, meaning that they always started matching at the start of the string. As of 1.2.0.pre46, regular expressions are not anchored, meaning that they can start their match anywhere in the string.

**4.7.6.8.5 IN** The IN comparison operator must have:

1. either an enumeration feature on the left hand side and a comma-separated list of enumeration constants in parentheses on the right-hand-side (or an object reference usage resolving to a list-of-enum-constants of the same type),
2. or an INTEGER feature on the left hand side, and a list of integers on the right-hand-side (or an object reference usage resolving to this),
3. or an ID\_D feature on the left hand side, and a list of integers on the right-hand-side (or an object reference usage resolving to a list of ID\_Ds).

For the first case, all of the enumeration constants must belong to the enumeration of the feature. The meaning is “feature must be in this list”, and is equivalent to a string of “=” comparisons with “OR” in between, and with parentheses around the string of OR-separated comparisons.

For the second and third cases, the meaning is the same, but applied to integers and `id_ds` respectively.

**4.7.6.8.6 HAS** The HAS comparison operator must have:

1. Either a list-of-enumeration constant on the left hand side and an enumeration constant belonging to the same enumeration on the left hand side (or an object reference usage resolving to this),
2. or a list-of-INTEGER feature on the left hand side, and an atomic integer value the right-hand-side (or an object reference usage resolving to this),
3. or a list-of-ID\_D feature on the left hand side, and an atomic `id_d` value on the right-hand-side (or an object reference usage resolving to this).

This signifies list-membership of the right-hand-side in the list on the left-hand-side.

## 4.7.7 Object references

### 4.7.7.1 Introduction

Object references are a way of referring to objects in a query outside of the object block which they matched. This provides the possibility of matching objects on the basis of the features of other objects earlier in the query.

In this subsection, we first give the syntax of object references, their declaration and their usage. We then provide some examples, followed by an explanation of those examples. We then give some explanation of object references. Finally, we document some constraints that exist on object references.

### 4.7.7.2 Syntax

```
object_reference_declaration : "AS" object_reference
;
object_reference : T_IDENTIFIER
;
object_reference_usage : object_reference
 "." feature_name
;
feature_name : T_IDENTIFIER
;
```

### 4.7.7.3 Examples

1. [Clause  
 [Phrase as p  
 phrase = NP  
 ]  
 ..  
 [Phrase  
 phrase = AP  
 and case = p.case  
 and number = p.number  
 and gender = p.gender  
 ]  
]
2. [Clause as C  
 [Phrase  
 phrase\_type = NP  
 parent = C.self  
 ]  
]
3. [Sentence as S]  
 ..  
 [Sentence  
 head = S.self  
 ]  
]

### 4.7.7.4 Explanation of examples

Example 1 finds, within a clause, first an NP, followed by an arbitrary stretch of text, followed by an AP. The AP's case, number, and gender-features must be the same as the NP's case, number, and gender-features

respectively.

Example 2 finds a clause, and within the clause an NP which is a direct constituent of the clause. That is, its parent feature is an `id_d` which points to its parent in the tree. This `id_d` must be the same as the clause's "self" feature. See section 2.7.6 on page 21 for more information about the "self" feature.

Example 3 finds a sentence and calls it S. Then follows an arbitrary stretch of text. Then follows another sentence whose head feature is an `id_d` which points to the first sentence. That is, the second sentence is dependent upon the first sentence.

#### 4.7.7.5 Explanation

The `object_reference_declaration` non-terminal is invoked from the object blocks, right after the object type name. That is, the `object_reference_declaration` must be the first item after the object type name, if it is to be there at all, as in all of the examples above. The object reference declaration says that the object that matched this object block must be called whatever the `object_reference` is (e.g., "p", "C", and "S" in the examples above). Then object blocks later in the query can refer to this object's features on the right-hand-side of feature comparisons. See section 4.7.6 on page 84 for details of feature-comparisons.

The `object_reference_usage` non-terminal shows up on the right-hand-side of feature comparisons, as in the examples above. It consists of an object reference followed by a dot followed by a feature name.

#### 4.7.7.6 Constraints on object references

The following are the constraints on object references:

- Object references must be declared before they can be used. That is, they must appear in an `object_reference_declaration` earlier in the query (i.e., further towards the top).
- The feature name on an object reference usage must be a feature of the object type of the object that had the corresponding object reference declaration.
- The feature type of the object reference usage must be the same as the feature type of the feature with which it is compared (not just compatible with).
- An object reference must only be declared once in a query. That is, no two object references must have the same name.
- A "Kleene Star" construct (see Section 4.7.8.3 on the following page) acts as an "export barrier" upwards in the tree for object reference declarations. Thus any object reference usages which are separated from the object reference declaration by a Kleene Star cannot be "seen". For example, this is not allowed:

```
[Clause
 [Phrase
 [Word as w1]
]* // Kleene Star acts as an export barrier!
 [Word surface=w1.surface] // So we can't see the declaration here...
]
```

- You also cannot have an object reference declaration on an object block that itself bears the Kleene Star. Thus this is not allowed:

```
[Clause
 [Phrase as p1]* // This is NOT allowed!
 [Phrase function=p1.function]
]
```

- It is not allowed to have an object reference declaration that is used “above” an OR. That is, all object reference declarations and usages should be within the same `block_str` (see Section 4.8.4 on page 94 and Section 4.8.5). “OR” acts as an “export barrier” on object reference declarations, not allowing them to be “seen” beyond the “OR”. Thus this is *not* allowed:

```
[Phrase as p1]
OR
[Phrase function=p1.function] // Oops! Can't see declaration from here!
Whereas this is allowed:
[Phrase as p1
 [Phrase function=p1.function] // This is OK
 OR
 [Phrase function<>p1.function] // This is also OK
]
```

The reason the second is allowed but the first is not is that it is the object reference *declaration* which is under embargo above (not below) an OR, whereas the object reference *usage* is free to see an object reference *declaration* that has been declared above an OR.

## 4.7.8 Block

### 4.7.8.1 Introduction

The non-terminal block is a choice between three kinds of block: `opt_gap_blocks`, `gap_blocks`, and `object_blocks`. It is used in the grammar of MQL queries in the definition of `block_strs`, that is, in when defining strings of blocks. See section 4.8.5 on page 95 for more information on `block_strs`.

### 4.7.8.2 Syntax

```
block : opt_gap_block
 | gap_block
 | object_block
 | ("NOTEXIST" | "NOTEXISTS") object_block
 | object_block "*" [monads_set]
;

```

### 4.7.8.3 The “\*” construct

The “\*” construct is a so-called “Kleene Star”. It allows searching for object blocks that are repeated. It has two forms: One with and one without a trailing set of integers (with the same syntax as a set of monads). For example:

```
SELECT ALL OBJECTS
WHERE
[Sentence
 [Phrase FIRST phrase_type = NP]
 [Phrase phrase_type IN (VP, NP, AP)]*
 [Phrase function = adjunct]* {1-3}
]
GO
```

This finds sentences whose first phrase is an NP, followed by arbitrarily many phrases which can either be VPs, NPs, or APs (or any combination of those), followed by between 1 and 3 phrases whose function is adjunct.

A less contrived example:

```

SELECT ALL OBJECTS
WHERE
[Sentence
 [Word psp=verb]
 [Word psp=article or psp=noun
 or psp=adjective or psp=conjunction]*{1-5}
]
GO

```

This finds sentences where there exists a word whose part of speech is verb, followed by 1 to 5 words whose parts of speech may be article, noun, adjective, or conjunction (or any combination of those). Presumably this would (in English) be a VP with (parts of) the object NP after the verb.

The Kleene-Star without a trailing set of integers means "from 0 to MAX\_MONADS". Note, however, that there is no performance penalty involved in such a large end: The algorithm stops looking when getting one more fails.

If 0 is in the set of integers, then this means that the object need not be there. This means that the following:

```

SELECT ALL OBJECTS
WHERE
[Sentence
 [Word psp=verb]
 [Word psp=article]*{0-1}
 [Word psp=noun]
]
GO

```

would find sentences where there exists a verb followed by 0 or 1 articles followed by a noun. Thus the case of "verb immediately followed by noun" would also be found by this query. Thus the "{0-1}" is equivalent to "?" in regular expressions.

The set of integers has the same syntax as monad sets. Therefore, to obtain a "no upper bound" star-construct, use {<lower-bound>-}, e.g., {10-} to mean "from 10 times to (practically) infinitely many times."

The following restrictions apply:

- You cannot have a Kleene Star on an object block which also has the NOTEXIST keyword in front.
- You cannot have a Kleene Star on an object block which has the "noretrieve" keyword.

## 4.8 Strings of blocks

### 4.8.1 Introduction

Having now described all the syntax and semantics of individual blocks, we now go on to giving the bigger picture of MQL queries. This section describes strings of blocks, as well as the higher-level non-terminals in the MQL query-query subset.

We first describe the `topograph`, the top-level entry-point into the MQL query-query grammar (4.8.2). We then describe the `blocks` non-terminal, which shows up inside each of the three kinds of blocks as an inner `blocks` (4.8.3). We then describe the `block_str` non-terminal, which provides for strings of blocks optionally connected by power blocks (the ". ." blocks which have been exemplified previously, and which mean "an arbitrary stretch of space") (4.8.5).

## 4.8.2 topograph

### 4.8.2.1 Introduction

The topograph non-terminal is the entry-point for the MQL query-query subset.<sup>2</sup> It simply consists of a blocks non-terminal. The topograph passes on a Universe and a Substrate to the blocks non-terminal, and these will be described below.

### 4.8.2.2 Syntax

```
topograph : blocks
;
```

### 4.8.2.3 Examples

1. [Word]
2. [Word psp=article]  
[Word psp=noun]
3. [Clause  
[Phrase phrase\_type = NP]  
..  
[Phrase phrase\_type = VP]  
]
4. [Book  
title = "Moby Dick"  
[Chapter chapter\_no = 3  
[Paragraph  
[Word surface = "Ishmael"]  
]  
..  
[Paragraph  
[Word surface = "whaling"]  
]  
]  
]

### 4.8.2.4 Explanation of examples

Example 1 simply finds all words within the topograph's Universe and Substrate.

Example 2 finds all pairs of adjacent words in which the first word is an article and the second word is a noun, within the topograph's Universe and Substrate.

Example 3 finds all clauses within which there are pairs of first an NP, followed by an arbitrary stretch of monads, then a VP. Within the topograph's Universe and Substrate, of course.

Example 4 finds a book whose title is "Moby Dick", and within the book it finds chapter 3, and within this chapter it finds a Paragraph within which there is a word whose surface is "Ishmael". Then, still within the chapter, after an arbitrary stretch of monads, it finds a Paragraph within which there is a word whose surface is "whaling".

---

<sup>2</sup>Even though `mql_query` is really the proper entry-point for an MQL query-query, we may consider the topograph to be the top-level syntactic non-terminal in the MQL query-subset. The topograph has historical primacy, since it was defined first in Doedens' QL (see [Doedens94]). The `mql_query` non-terminal simply acts as a proxy, passing control to the topograph immediately.

### 4.8.2.5 Universe and Substrate

In order to understand how the Universe and Substrate are calculated, it is necessary to refer back to the definition of the SELECT OBJECTS query. Please consult section 3.8.1.4 on page 45 for details.

## 4.8.3 blocks

### 4.8.3.1 Introduction

The `blocks` non-terminal is central in the MQL query-query subset. It shows up in five places:

- In the `topograph`,
- Inside the `object_block` as the inner `blocks`,
- Inside the `gap_block` and the `opt_gap_block` as the inner `blocks`.

### 4.8.3.2 Syntax

```
blocks : block_string
;
```

## 4.8.4 block\_string

### 4.8.4.1 Introduction

A “`block_string`” is basically either a “`block_str`” or it is a `block_str` followed by the keyword “OR” followed by another `block_string`. Or, put another way, a `block_string` is a string (possibly 1 long) or `block_str`’s, separated by “OR”.

### 4.8.4.2 Syntax

```
block_string : block_str
 | block_str "OR" block_string
;
```

### 4.8.4.3 Examples

1. [Clause  
    [Phrase function = Predicate] // This...  
    [Phrase function = Objc] // ... is a block\_str  
    OR  
    [Phrase function = Predicate] // And this...  
    [Phrase function = Complement] // is another block\_str  
  ]
2. [Sentence  
    [gap [Clause function = relative] // This is a block\_str  
    OR  
    [Clause AS c1 function = Subject] // And this...  
    .. // ... is also ...  
    [Clause daughter = c1.self] // ... a block\_str  
  ]

#### 4.8.4.4 Explanation

The OR construct between strings of blocks works like this: First, the sheaf for the first `block_str` is found. Then the sheaf for the `block_string` is found recursively. When it returns, do the following: If both sheaves are failed sheaves, return a failed sheaf. Otherwise, if only one of them is a failed sheaf, return the sheaf that is not a fail. If both of them are non-failed sheaves, take the sheaf from the `block_string` and append it to the sheaf from the `block_str`.

The first example finds all clauses in which it is either the case that there exist two phrases inside the clause where the first is a predicate and right next to it is an object, or the first is a predicate and right next to it is a complement (or both might be true, in which case you'll get two straws inside the inner sheaf of the clause).

The second example finds all clauses in which it is the case that there either is a gap with a relative clause inside it, or there are two clauses (possibly separated) where the first clause is a subject in the second. (This assumes a data model where mother clauses do not include the monads of their daughter clauses).

See Section 4.7.7 on page 89 for some restrictions on object references regarding OR.

### 4.8.5 `block_str`

#### 4.8.5.1 Introduction

The `block_str` non-terminal shows up as the only non-terminal in the rule for the `blocks` non-terminal (see the previous section). Thus it defines what a `blocks` is.

A `block_str` is simply a string of `block_strs`. If there is more than one `block_str`, then they must be separated by a `power` block. A `power` block looks like “`..`” or “`.. < 5`” or “`.. BETWEEN 1 AND 5`”. If it is “`..`”, it means “here must come a stretch of monads of arbitrary length, which could be no monads (0 length)”. Otherwise, there is an upper and/or lower limit on the number of monads that can intervene. The `block_str`, in turn, is a string of blocks of the four kinds described in section 4.7.

In the following, we first define the syntax relevant to `block_strs`. We then give a number of examples, after which we describe the `power` block. Finally, we described the “bang” operator, or “`!`”.

#### 4.8.5.2 Syntax

```
block_str : block
 | block "!" block_str
 | block block_str
 | block power block_str
;
power : ".." [restrictor]
;
restrictor : "<" limit
 | "<=" limit
 | "BETWEEN" limit "AND" limit
;
limit : T_INTEGER /* non-negative integer, may be 0. */
;
```

#### 4.8.5.3 Examples

1. [Word]
2. [Word psp=article]  
[Word psp=noun]  
.. <= 5  
[Word psp=verb]

```

3. [Phrase phrase_type = NP]
 ..
 [Phrase phrase_type = AdvP]
 .. BETWEEN 1 AND 5
 [Phrase phrase_type = VP]

4. [Chapter
 topic = "Noun classes in Bantu"
]
 [Chapter
 topic = "Causatives in Setswana"
]
 ..
 [Chapter
 topic = "Verb-forms in Sesotho"
]

```

#### 4.8.5.4 power

The power block means “before the start of the next block, there must come a stretch of monads of arbitrary length, which can also be no monads (0 length)”. In its basic form, it is simply two dots, “..”.

The stretch of monads is calculated from the monad after the last monad of the previous block. If the previous block ended at monad 7, then the power block starts counting monads from monad 8.

One can optionally place a *restrictor* after the two dots, thus making the power block look like this, e.g., “.. < 5”, “.. <= 5”, or “.. BETWEEN 1 AND 5”.

The first two kinds of restrictor mean “although the stretch of monads is of arbitrary length, the length must be less than (or equal to) the number of monads given in the restrictor”. Thus “.. < 5” means “from 0 to 4 monads after the end of the previous block”, and “.. <= 5” means “from 0 to 5 monads after the end of the previous block”. That is, if the previous block ended at monad 7, then “.. < 5” means “the next block must start within the monads 8 to 12”, while “.. <= 5” means “the next block must start within the monads 8 to 13”.

Similarly, the third kind, “.. BETWEEN *min* AND *max*” means “there must be at least *min* monads in between, and at most *max* monads. This is construed as “>= *min* AND <= *max*”.

#### 4.8.5.5 The bang (“!”)

You can optionally place a bang (“!”) between any of the blocks in a *block\_str*. The bang indicates that there must be no gaps between the two blocks. It is an implicit rule of MQL that there is a hidden *opt\_gap\_block* between each pair of blocks in a *block\_str* which are not mediated by a bang.

The reason for having the *opt\_gap\_block* is the following: It protects you from what you do not know. In some languages, there can be gaps in clauses because of post-positive conjunctions “sticking out” at a higher level. One would not wish to have to specify all the time that one wanted to look for gaps, for one would invariably forget it sometimes, thus not getting all the results available. Thus MQL inserts an *opt\_gap\_block* between each pair of blocks that are not mediated by a bang. The bang is a way of specifying that one does not wish the hidden *opt\_gap\_block* to be inserted.

The *opt\_gap\_block* that is inserted is not retrieved.

# Appendix A

## Copying

### A.1 Introduction

Emdros is covered by two different licenses, both of which allow you freely to copy, use, and modify the sourcecode. The parts which were written by Ulrik Petersen are covered by the GNU GPL. The pcre library, which provides regular expressions-capabilities, is covered by a different license. Some parts were contributed by Kirk E. Lowery or Chris Wilson; they are Copyright Ulrik Petersen and are also under the GNU GPL.

SQLite is in the Public Domain. See [www.sqlite.org](http://www.sqlite.org) for details.

All Emdros documentation (including this document) is covered under the Creative Commons Attribution-Sharealike license version 2.5. This license is included below.

### A.2 Creative Commons Deed (for all documentation)

Attribution-ShareAlike 2.5

CREATIVE COMMONS CORPORATION IS NOT A LAW FIRM AND DOES NOT PROVIDE LEGAL SERVICES. DISTRIBUTION OF THIS LICENSE DOES NOT CREATE AN ATTORNEY-CLIENT RELATIONSHIP. CREATIVE COMMONS PROVIDES THIS INFORMATION ON AN "AS-IS" BASIS. CREATIVE COMMONS MAKES NO WARRANTIES REGARDING THE INFORMATION PROVIDED, AND DISCLAIMS LIABILITY FOR DAMAGES RESULTING FROM ITS USE.

#### *License*

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED.

BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

#### 1. Definitions

- (a) "**Collective Work**" means a work, such as a periodical issue, anthology or encyclopedia, in which the Work in its entirety in unmodified form, along with a number of other contributions, constituting separate and independent works in themselves, are assembled into a collective whole. A work that constitutes a Collective Work will not be considered a Derivative Work (as defined below) for the purposes of this License.

- (b) **"Derivative Work"** means a work based upon the Work or upon the Work and other pre-existing works, such as a translation, musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which the Work may be recast, transformed, or adapted, except that a work that constitutes a Collective Work will not be considered a Derivative Work for the purpose of this License. For the avoidance of doubt, where the Work is a musical composition or sound recording, the synchronization of the Work in timed-relation with a moving image ("synching") will be considered a Derivative Work for the purpose of this License.
  - (c) **"Licensor"** means the individual or entity that offers the Work under the terms of this License.
  - (d) **"Original Author"** means the individual or entity who created the Work.
  - (e) **"Work"** means the copyrightable work of authorship offered under the terms of this License.
  - (f) **"You"** means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.
  - (g) **"License Elements"** means the following high-level license attributes as selected by Licensor and indicated in the title of this License: Attribution, ShareAlike.
2. **Fair Use Rights.** Nothing in this license is intended to reduce, limit, or restrict any rights arising from fair use, first sale or other limitations on the exclusive rights of the copyright owner under copyright law or other applicable laws.
3. **License Grant.** Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:
- (a) to reproduce the Work, to incorporate the Work into one or more Collective Works, and to reproduce the Work as incorporated in the Collective Works;
  - (b) to create and reproduce Derivative Works;
  - (c) to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission the Work including as incorporated in Collective Works;
  - (d) to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission Derivative Works.
  - (e) For the avoidance of doubt, where the work is a musical composition:
    - i. **Performance Royalties Under Blanket Licenses.** Licensor waives the exclusive right to collect, whether individually or via a performance rights society (e.g. ASCAP, BMI, SESAC), royalties for the public performance or public digital performance (e.g. webcast) of the Work.
    - ii. **Mechanical Rights and Statutory Royalties.** Licensor waives the exclusive right to collect, whether individually or via a music rights society or designated agent (e.g. Harry Fox Agency), royalties for any phonorecord You create from the Work ("cover version") and distribute, subject to the compulsory license created by 17 USC Section 115 of the US Copyright Act (or the equivalent in other jurisdictions).
  - (f) **Webcasting Rights and Statutory Royalties.** For the avoidance of doubt, where the Work is a sound recording, Licensor waives the exclusive right to collect, whether individually or via a performance-rights society (e.g. SoundExchange), royalties for the public digital performance (e.g. webcast) of the Work, subject to the compulsory license created by 17 USC Section 114 of the US Copyright Act (or the equivalent in other jurisdictions).

The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats. All rights not expressly granted by Licensor are hereby reserved.

4. **Restrictions.** The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

- (a) You may distribute, publicly display, publicly perform, or publicly digitally perform the Work only under the terms of this License, and You must include a copy of, or the Uniform Resource Identifier for, this License with every copy or phonorecord of the Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Work that alter or restrict the terms of this License or the recipients' exercise of the rights granted hereunder. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties. You may not distribute, publicly display, publicly perform, or publicly digitally perform the Work with any technological measures that control access or use of the Work in a manner inconsistent with the terms of this License Agreement. The above applies to the Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Work itself to be made subject to the terms of this License. If You create a Collective Work, upon notice from any Licensor You must, to the extent practicable, remove from the Collective Work any credit as required by clause 4(c), as requested. If You create a Derivative Work, upon notice from any Licensor You must, to the extent practicable, remove from the Derivative Work any credit as required by clause 4(c), as requested.
- (b) You may distribute, publicly display, publicly perform, or publicly digitally perform a Derivative Work only under the terms of this License, a later version of this License with the same License Elements as this License, or a Creative Commons iCommons license that contains the same License Elements as this License (e.g. Attribution-ShareAlike 2.5 Japan). You must include a copy of, or the Uniform Resource Identifier for, this License or other license specified in the previous sentence with every copy or phonorecord of each Derivative Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Derivative Works that alter or restrict the terms of this License or the recipients' exercise of the rights granted hereunder, and You must keep intact all notices that refer to this License and to the disclaimer of warranties. You may not distribute, publicly display, publicly perform, or publicly digitally perform the Derivative Work with any technological measures that control access or use of the Work in a manner inconsistent with the terms of this License Agreement. The above applies to the Derivative Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Derivative Work itself to be made subject to the terms of this License.
- (c) If you distribute, publicly display, publicly perform, or publicly digitally perform the Work or any Derivative Works or Collective Works, You must keep intact all copyright notices for the Work and provide, reasonable to the medium or means You are utilizing: (i) the name of the Original Author (or pseudonym, if applicable) if supplied, and/or (ii) if the Original Author and/or Licensor designate another party or parties (e.g. a sponsor institute, publishing entity, journal) for attribution in Licensor's copyright notice, terms of service or by other reasonable means, the name of such party or parties; the title of the Work if supplied; to the extent reasonably practicable, the Uniform Resource Identifier, if any, that Licensor specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work; and in the case of a Derivative Work, a credit identifying the use of the Work in the Derivative Work (e.g., "French translation of the Work by Original Author," or "Screenplay based on original Work by Original Author"). Such credit may be implemented in any reasonable manner; provided, however, that in the case of a Derivative Work or Collective Work, at a minimum such credit will appear where any other comparable authorship credit appears and in a manner at least as prominent as such other comparable authorship credit.

5. **Representations, Warranties and Disclaimer**

UNLESS OTHERWISE AGREED TO BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND

CONCERNING THE MATERIALS, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

**6. Limitation on Liability.**

EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

**7. Termination**

- (a) This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Derivative Works or Collective Works from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.
- (b) Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

**8. Miscellaneous**

- (a) Each time You distribute or publicly digitally perform the Work or a Collective Work, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.
- (b) Each time You distribute or publicly digitally perform a Derivative Work, Licensor offers to the recipient a license to the original Work on the same terms and conditions as the license granted to You under this License.
- (c) If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.
- (d) No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.
- (e) This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.

Creative Commons is not a party to this License, and makes no warranty whatsoever in connection with the Work. Creative Commons will not be liable to You or any party on any legal theory for any damages whatsoever,

including without limitation any general, special, incidental or consequential damages arising in connection to this license. Notwithstanding the foregoing two (2) sentences, if Creative Commons has expressly identified itself as the Licensor hereunder, it shall have all rights and obligations of Licensor.

Except for the limited purpose of indicating to the public that the Work is licensed under the CCPL, neither party will use the trademark "Creative Commons" or any related trademark or logo of Creative Commons without the prior written consent of Creative Commons. Any permitted use will be in compliance with Creative Commons' then-current trademark usage guidelines, as may be published on its website or otherwise made available upon request from time to time.

Creative Commons may be contacted at <http://creativecommons.org/>.

### **A.3 GNU General Public License**

General Public License (GPL)

GNU GENERAL PUBLIC LICENSE  
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the

source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

#### GNU GENERAL PUBLIC LICENSE

##### TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you". Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it. Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to

be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### END OF TERMS AND CONDITIONS

#### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
This program is free software; you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation; either version 2 of the License, or
(at your option) any later version.
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
```

Also add information on how to contact you by electronic and paper mail. If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program
`Gnomovision' (which makes passes at compilers) written by James Hacker.
<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

## A.4 PCRE license

PCRE is a library of functions to support regular expressions whose syntax and semantics are as close as possible to those of the Perl 5 language.

Release 6 of PCRE is distributed under the terms of the "BSD" licence, as specified below. The documentation for PCRE, supplied in the "doc" directory, is distributed under the same terms as the software itself.

The basic library functions are written in C and are freestanding. Also included in the distribution is a set of C++ wrapper functions.

### A.4.1 THE BASIC LIBRARY FUNCTIONS

Written by: Philip Hazel

Email local part: ph10

Email domain: cam.ac.uk

University of Cambridge Computing Service,

Cambridge, England. Phone: +44 1223 334714.

Copyright (c) 1997-2006 University of Cambridge

All rights reserved.

## **A.4.2 THE C++ WRAPPER FUNCTIONS**

Contributed by: Google Inc.

Copyright (c) 2006, Google Inc. All rights reserved.

## **A.4.3 The “BSD” license**

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the University of Cambridge nor the name of Google Inc. nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Appendix B

# Console sheaf grammar

### B.1 Introduction

The sheaf's contents were explained in section 4.4 on page 75. In this appendix, we give the grammar for the sheaf as it is output with console output (as opposed to XML output).

### B.2 Sheaf grammar

```
/* Sheaf */
sheaf : failed_sheaf | successful_sheaf
;
failed_sheaf : "//" /* A failed sheaf means
 that the query failed
 in some way. */
;
successful_sheaf : "//" straws /* A successful sheaf
 means that the query
 did not fail. It may
 however, be empty,
 in which case the
 list_of_straws will
 not be there, and the
 sheaf will look like
 this: "// < >".
 */
;
straws : "<" list_of_straws ">"
;
list_of_straws : { straw }
;

/* Straw */
straw : "<" list_of_matched_objects ">"
;
list_of_matched_objects : { matched_object }
;
/* Matched object */
matched_object : mo_id_d | mo_id_m
;
```

```

/* Matched object with id_d */
mo_id_d : "[" object_type_name
 id_d_of_object
 monad_set
 is_focus
 inner_sheaf
 "]"
;
object_type_name : T_IDENTIFIER
;
id_d_of_object : T_INTEGER
;
is_focus : "true" | "false" /* Was the block against
 which this matched_object
 was matched a "focus"
 block or not? I.e., was
 they keyword "focus"
 present in the block?
 */
;
inner_sheaf : sheaf
;

/* Matched object with id_m (see [Standard-MdF]) */
mo_id_d : "[" "pow_m"
 monad_set
 is_focus
 inner_sheaf
 "]"
;

```

### B.3 References

For the `monad_set` non-terminal, please see section 3.8.1 on page 43.

# Bibliography

- [Doedens94] Doedens, Crist-Jan. 1994: *Text Databases, One Database Model and Several Retrieval Languages*. 'Language and Computers,' Volume 14. Editions Rodopi Amsterdam, Atlanta, GA. ISBN 90-5183-729-1.
- [Standard-MdF] Petersen, Ulrik. 2002: *The Standard MdF Model*. Unpublished article. Obtainable from URL: <http://emdros.org/>
- [Relational-EMdF] Petersen, Ulrik. 2005: *Relational Implementation of EMdF and MQL*. Unpublished working-paper. Obtainable from URL: <http://emdros.org/>
- [MQLQueryGuide] Petersen, Ulrik. 2005: *MQL Query Guide*. Obtainable from URL: <http://emdros.org/>
- [Monad Sets] Petersen, Ulrik. 2002: *Monad Sets – Implementation and Mathematical Foundations*. Unpublished article. Obtainable from URL: <http://emdros.org/>

# Index

- all\_m, 19, **20**, 37, 44, 45, 58, 59, 61, 62
- any\_m, 19, **20**, 37, 58, 59, 61, 62
- ASCII, 18, 50
  
- Backus-Naur Form, 10, **11**, 13
- borders, **21**, 54, 80
  
- case-sensitivity, 25
- computed feature, 37, 49
- console
  - output, 26, 27, 45, 108
  - sheaf, 45, 108
  
- Doedens, Crist-Jan, 10, 68, 93, 110
  
- EMdF, 10
  - acronym, 10
  - database, 10, 16–20
    - example, 19
  - database engine, *see* Emdros, 10, 16
  - model, 10, 16, **16**, 17–19
  - origins, 10
- Emdros
  - license, 101
  - origins, 10
- enumeration, 18, **21**, 30, 37, 50, 88
  - constant, 22, 26, 37, 40, 42, 51, 57, 74, 85, 87, 88
    - comparison, 88
    - equality, 88
    - querying, 42, 50, 88
    - retrieving, 42
  - creation, 40
  - database consistency, 42
  - default constant, **22**, 40, 42
  - deletion, 42
  - dropping, 42
  - manipulation, 40
  - namespace, 22, 26
  - object types using, 42, 51
  - querying, 50
  - update, 41
  
- feature
  - computed, *see* computed feature
  - example, 19
  
- flat sheaf, 45, **77**
- flatten operator, 45, **77**
- full access language, 10
  
- gaps, **21**, 54, 71, 76–78, 80–83, 96
  
- Hazel, Philip, 15
  
- id\_d, 18, **20**, 50
- id\_m, **20**
- identifier, 24
  - case-sensitivity of, *see* case-sensitivity
- integer, 18, 50
  
- max\_m, 20, **22**, 52
- MdF, 10
  - origins, 10
- min\_m, 20, **22**, 51
- monad
  - example, 19
- MQL, 10
  - acronym, 10
  - origins, 10
- mql(1) program, 26, 67
  
- Namespace, **25**
- namespace, 22
  
- object, 30
  - example, 19
  - id, 18, **20**
  - id\_d, **20**
  - id\_m, **20**
- object type, 30
  - example, 19
- object types
  - special, *see* all\_m, any\_m, pow\_m
- output
  - console, 26, 27, 45, 108
  - XML, 26, 27, 108
  
- part\_of, 20, 21, 45, 61, 63, 81
- PCRE
  - library, **15**, 106
  - license, 106
- Perl 5, 88

pow\_m, 19, 20, **20**, 21, 37, 58, 59, 61, 62

QL, 10

regular expressions, **88**

  match, 86

  support, 10, 15

  syntax, 74, 85, 87

self, **21**, 37, 49, 61, 62

  creation of, 37

  example using, 89

  type of, 21

sheaf, 26, 45, 69, **75**, 76, 83, 108

  flat, *see* flat sheaf

string, 18, 50

Su, *see* substrate

substrate, 45, 71, 76, 77, **77**, 80–82, 84, 93, 94

T\_IDENTIFIER, 14, 24, **24**, 25, 28–32, 35, 36, 38–  
43, 46, 48–51, 53, 55, 57, 61, 62, 73, 74,  
79, 85, 89, 109

T\_INTEGER, 14, 24, **25**, 36, 43, 44, 56, 58, 74, 95,  
109

T\_STRING, 24, 25, **25**, 28–30, 36, 57, 74, 85

types, 18

  7-bit string, 18

  8-bit string, 18

  ascii, 18

  compatibility, **87**

  enumeration, 18

  id\_d, 18

  integer, 18

  string, 18

U, *see* universe

universe, 45, 70, 77, **77**, 80, 82, 84, 93, 94

XML

  output, 26, 27, 108