

Update Plans

A High Level Low Level Specification Language

Hugh Osborne

Department of Informatics
Faculty of Mathematics and
Informatics
University of Nijmegen
Toernooiveld
6525 ED Nijmegen
The Netherlands

Table of Contents

Table of Contents	iii
Preface	vii
Chapter One: Introduction	1 Update Plans 2 2 Related Approaches 4 3 Some Examples 5 4 Overview 7 5 Notational Conventions 7
Chapter Two: Syntax	1 Basic Syntax 9 2 Sweetening the Syntax 10 3 Syntactic Sugar 11 4 Notational Conventions 12 5 Further Simplifications 16 6 Summary 17
Chapter Three: Semantics	1 Memory 19 2 Update Schemes 21

Chapter Four: An Introductory Application and Example	1 Compiling λ MMachine Programmes to FLIP code 26 2 The Basic λ -Calculus Machine 28 3 Basic Values 31 4 Strict Abstraction 32 5 Predefined Functions 35 6 I/O 36 7 The Implementation 37
Chapter Five: Typing	1 Introduction 39 2 Typing 40 3 Ground Expressions 41 4 Memory Structures 43 5 Programme Stores 46 6 Conclusions 46
Chapter Six: Semantic Properties of Update Plans	1 Semantic Equivalence 47 2 Box Diagrams 49 3 Restricted and General Properties 50 4 Semantic Irrelevance 50 5 Deriving Transitive from Local Properties 53 6 Conclusions 54
Chapter Seven: Archetypes	1 Archetypes 55 2 Syntactic Sugar 59 3 Expansion 60 4 Recursion 64 5 Some Examples 67 6 Conclusions 69
Chapter Eight: Demonstrably Correct Linearisation of Intermediate Code	1 The Linear Code Machine 72 2 The Tree Machine 76 3 Linearising Tree Code 79 4 Proving Semantic Equivalence 81 5 Conclusions 85
Chapter Nine: Synchronous Parallelism	1 Informal Introduction 87 2 Syntax 87 3 Formal Specification 88 4 An Example: The Berkeley RISC II CPU 90 5 Conclusions 94

Chapter Ten: Implementational Aspects	1 Notes on Notation 95 2 Basic Update Schemes 97 3 Archetypes 102 4 Backtracking 108 5 Conclusions 112
Chapter Eleven: Conclusions and Suggestions for Further Research	1 Conclusions 115 2 Further Research 115
Summary	119
Samenvatting	121
Appendix I: Bibliography	I.1
Appendix II: Grammar	II.1
Appendix III: The FLIP Machine	III.1
Appendix IV: Store Structures	IV.1
Appendix V: Semantic Properties	V.1

**Appendix VI:
The Linear and
Tree Machines**

VI.1

**Appendix VII:
The Berkeley
RISC II processor**

VII.1

**Appendix VIII:
Glossary**

VIII.1

This thesis is the result of work carried out between 1989 and 1994. During the first two years of this period I was a research assistant, at the University of Nijmegen, in Esprit Basic Research Action no. 3147 (the Phoenix project), carried out in collaboration with Imperial College in London and GMD in Karlsruhe, and financed by the European Community. The remainder of the period I was an “AIO” (postgraduate research assistant) at the University of Nijmegen. I am grateful to both the European Community and the University of Nijmegen (and thus, indirectly, to the Dutch and European tax payers) for financing this research.

Some of the work in this thesis has been presented earlier. In particular chapters two, three and four are based on work presented in [56, 57, 58, 59, 60].

I would like to thank the many people who have contributed to the development of the ideas presented in this thesis. First and foremost, I would like to thank Hans Meijer, who originally developed Update Plans (then called Update Schemes) [53], and who has played a fundamental rôle in helping to develop and formalise the key concepts in this thesis. I would also like to thank Robert Giegerich, of the University of Bielefeld, whose ideas formed the inspiration for chapter eight and, indirectly, chapters seven and six. I would, of course, also like to thank my colleagues on the Phoenix project, and in particular Hendrick Lock (GMD, Karlsruhe), whose JUMP machine [47, 48] formed one of the first test cases for Update Plans, and Erik Meijer (KUN, Nijmegen) and Ross Patterson (Imperial, London), whose λ MMachine [52] is specified in chapter four. Other people I would like to thank are Mark van den Brand, Todd Cook, Max Geerling, Wil Lamain, Kees Koster, Ineke Kusters, Greta Löw and Harrie van Seters. The list would, of course, not be complete without mentioning Mariël, Sara and Tobias. Their patience and encouragement (“Papa, wanneer is jouw boekje nou af?” [66]) were essential to the completion of this work. Sara deserves a special mention for designing the cover.

It has been said [38] that computer programming is an art, and this is undoubtedly true — if the products of civil engineers had to be “debugged” as often as those of software engineers do, the world would be a hazardous place. Besides threatening to banish computer science to a no-man’s land between C.P. Snow’s two cultures [69], this lack of precision has more serious consequences. As computer systems become ubiquitous, the software engineer’s products begin to play as important a part in daily life as do the civil engineer’s. One of the major tasks, therefore, of computer science — perhaps the subject’s whole *raison d’être* — is to supply formalisms and methodologies for constructing provably correct software. Historically this problem has been addressed on three fronts.

Methodologies have been invented for the development of software by the transformation of precise but conceptually simple specifications of problems to less intuitive but more easily implementable programmes [5, 7, 8, 9, 50, 51, 62, 70, 71].

Programming languages have become more and more precisely defined [54, 75] making precise definitions of the semantics of programmes more and more feasible.

Compilers are of course themselves programmes. By a combination of these two techniques formal specifications of languages are transformed to working compilers [20, 25, 33, 34, 43, 53].

These three methods all address the problem of translating or compiling problem specifications in some high level language to an equivalent specification in a low level language, mainly by means of syntactic transformations. The semantics of the low level language are usually not (formally) specified, but assumed to be trivial. This is an unwarranted assumption.

This classic example, from the architecture of the PDP-11, shows the dangers of assuming the semantics of low level languages to be trivial. Consider the predecrement and postincrement addressing modes. The semantics seem simple. In the predecrement mode an address in a register is decreased by one (byte or word) and used as the address of the operand, e.g. if register `R6` contains the address `1234` then the object accessed by `-(R6)` is the object at address `1233`, and this is also the address that will be found in `R6` after the access. The postincrement mode is the inverse, first the object is accessed and then the address is increased. Addressing mode `(R6)+` will address the object at address `1234`, after which `R6` will contain `1235`.

A formal specification of the semantics of these addressing modes seems superfluous. In combination, however, they are less simple. Take, for example, the instruction `'MOV (R6)+ -(R6)'`. What does it mean? Is the postincrement executed, and then the predecrement, or is the order of access random? Maybe the arguments are accessed simultaneously, but what does this mean? Perhaps a formal specification is necessary after all.

Formal specifications, derived using the methods above, cannot be considered complete without attention to the details of the semantics of the underlying low level language — usually the instruction set of some machine, either concrete or abstract. A formalism is needed in which the semantics of such low level languages can be specified. Providing an important link in formal specification methods is not the only use of such a formalism. In the case of specifications of abstract machines, or concrete machines under development, the specification can be useful in reaching design decisions, and optimising the instruction set. In the case of existing concrete machines it can be used for deriving provably correct optimisations of generated code.

Update Plans are intended as a well defined specification formalism, with sufficient abstractive power to facilitate specification, yet close enough to concrete machine models to enable simple implementation of specification based prototypes.

1 Update Plans

Update Plans are a form of rewrite system, in particular a graph rewrite system [10, 21], closely related to term graph rewriting [6, 72] (though this will not be pursued in this thesis). They combine a well defined formal semantics, and an intuitive operational semantics with readability and flexibility. They were introduced by Meijer [53] to describe machine-code generated by a compiler in a machine independent way. They were deliberately designed in such a way that an update plan resembles a set of rewrite rules or, in a sugared version of the syntax, function definitions in

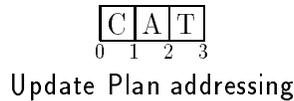
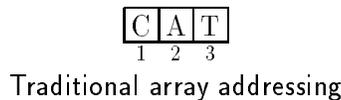


Figure 1: Addressing in arrays and update plans

a declarative language.

Update plans specify state transitions in some abstract machine. This machine consists of a number of *stores*, each containing a countable set of *cells* addressed by a completely ordered set of locations, called *locators*. Conceptually it is not the cells themselves that are addressed, but the boundaries between cells, so that a sequence is not considered to be *at* certain locators, but rather *between* locators. For example, rather than representing the string ‘CAT’ by an array of characters, and considering the substring ‘AT’ to be the subarray from locator 2 to locator 3, ‘CAT’ is considered to be a sequence of characters between locator 0 and locator 3, and ‘AT’ to be the subsequence between locators 1 and 3. This is illustrated in figure 1. The basic notation for such a sequence is $0[\text{‘CAT’}]3$. An extensive motivation for this numbering convention is given by Meijer [53]. Suffice it here to mention that $a[x\ y]c$ is the same as $a[x]b[y]c$, for some b , and that with $a[x]b$ we have $\text{length } x = b - a$.

The notation $a+n$ is used to refer to the locator n cells to the right of a , and $a-n$ to refer to the cell n cells to the left. A singleton sequence is identified with its element. Sequences are specified by triples of the form

$$\langle locator \rangle [\langle sequence \rangle] \langle locator \rangle ,$$

called *locator expressions*. A grammar for basic Update Plans is given in chapter two. A set of locator expressions is *consistent* if there are no two expressions in the set which specify the contents of some cell to have two different values. A consistent set of locator expressions describes a (sub)configuration of the machine.

The immediate constituent part of an update plan, an *update scheme*, is constructed from two sets of locator expressions forming the *left hand side* and the *right hand side*, and from a boolean expression, known as the *guard*. Variables (indicated by lower case words) are allowed in the constituent expressions of an update scheme. An update scheme containing only constants (indicated either by a value or, symbolically, by upper case words) is known as an *update rule*. Some simple examples can be found later in this chapter. Given a substitution and evaluation mechanism mapping variables to (sequences of) constants an update scheme can be instantiated to an update rule. Both the left and right hand sides of the resulting update rule must be self-consistent, i.e. all locator expressions in the left and right hand side respectively must be mutually consistent. The left and right hand side, however, need not be consistent with each other.

Briefly, Update Plans work as follows (more detailed descriptions are to be found in the remainder of this thesis). An update rule is *applicable* to a given configuration if its left hand side is a subset of that configuration and its guard is true. The result of applying an applicable update rule to a configuration c is the superset of the right hand side that is minimally different from c . A set of update schemes is called an *update plan*. One configuration can be *derived* from another (source) configuration by an update plan if the plan contains an update scheme an instantiation of which is applicable to the configuration. The derived configuration is then the result of applying such an instantiation to the source configuration. Configurations can also be derived by repeated applications of an update plan. A *final configuration* for an update plan is a configuration to which none of the update schemes in the plan is applicable. A *script* is an update plan together with an *initial configuration*. A *final development* of a script is a final configuration derived from the script's initial configuration by applications of the script's update plan.

A glossary of the terminology of Update Plans can be found in appendix VIII. A formal definition of the syntax is given in chapter two, and of the semantics in chapter three.

2 Related Approaches

Various other methods have been proposed for specifying low level activities. Abstract machines have been specified using transition systems [22], informal descriptions [73], an imperative programming style [47, 63] and functional languages [41], to name but a few. The best known contributions from the concrete side are probably ISPS [68] and register transfer languages, for example three address code [1]. These methods all have their drawbacks, lacking either a formal definition, or being impractical in that they are not easy to realise as, or translate to concrete implementations.

The drawbacks of informal specifications are well known — it could be said that ‘informal specification’ is an oxymoron, and that nothing better than a description can be given informally.

Transition systems have a well defined semantics, but quickly become unwieldy as the complexity of the system being specified increases. Even expressing a simple **JUMP** instruction leads, at best, to an inelegant use of the formalism. Functional languages also have a well defined formal semantics, but are also in general too far from concrete architectures to allow for intuitive specifications. Most current implementations of functional languages are also too inefficient for realistic prototyping. At the other end of the scale, a specification in some imperative language will be easy to implement with reasonable efficiency, but will suffer from a lack of a flexible and useful formal semantics. The same is true, possibly to a lesser extent, of ISPS and register transfer languages. Indeed, despite the existence of an ISP specification [19], early models of the PDP-11 had differing semantics for commands such as the **MOV (R6)+ -(R6)** from example 1, and the semantics of an instruction such as **ADD (R6)+ -(R6)** was not even unambiguously defined.

Most of these methods have been used as a back end for specification systems. The choice of back end is usually between a high level language, with the loss of efficiency and divergence from a final low level implementation that this entails; and a low level language, with the resulting loss of generality and portability. Update Plans are intended as a high level language for specifying low level activities, in such a way that these specifications are relatively simple to transform to specific low level languages. This is achieved by defining a declarative specification language — Update Plans — with an underlying imperative machine model. This makes update plans suitable both for specifying low level activities and as a back end for other specification formalisms. The *archetype* mechanism, which adds a macro-like mechanism to update plans, and which is introduced in chapter seven, possibly weakens the direct link between specification and implementation, but, as illustrated in chapter eight, provides a powerful mechanism for writing specifications which, by means of transformations within the same formalism, lead to easily implementable specifications.

At least two other specification languages combining imperative and declarative features have been developed [12, 13, 27, 29], but in these cases it is the underlying model which is declarative and the surface structure which has a strong imperative flavour. Reversing this, as in Update Plans, leads to a model in which it is easy to reason, and yet in which essentially imperative machine primitives can easily be expressed and combined. A possible application is an update plan based version of a peephole optimiser [15, 16, 26]. In the peephole optimiser ISP specifications of neighbouring commands are combined and an attempt is made to identify the resulting specification with some other command. The referential transparency of Update Plans is much more suited to such a process than imperative formalisms such as ISP and ISPS. The specification and proof of such an optimiser would make use of the techniques demonstrated in chapter eight.

Update Plans combine the facility of use of imperative languages with the formality of declarative languages, by providing a high level declarative language with an underlying low level imperative model.

3 Some Examples

The two scheme update plan in example 2(a) computes the greatest common divisor of the number initially between A and B and that initially between B and C.

————— Example 2 (a) —————

$$\begin{aligned} A[x]B B[y]C &= [x < y] \Rightarrow B[y - x]C. \\ A[x]B B[y]C &= [y < x] \Rightarrow A[x - y]B. \end{aligned}$$

Capitalised words denote (unspecified) constants: A, B and C are fixed locators, and x and y are variables. In fact, if at any stage of the computation we have A[9]B and B[6]C, (only) the update rule in example 2(b) is applicable, whereupon the 9 is replaced by a 3.

$$A[9]B \ B[6]C \Longrightarrow A[3]B.$$

The locators may likewise be specified by variables and obtain their actual value by instantiation, as illustrated by the (nondeterministic) update plan in example 3 which, in the context of a set of appropriate typing rules, sorts the sequence initially between **A** and **C**.

Example 3

$$A[xs]a \ a[x]b \ b[y]c \ c[ys]C \ = [x > y] \Rightarrow a[y]b \ b[x]c.$$

By simple notational conventions “irrelevant” addresses may be omitted and adjacent locator expressions combined (the concept “irrelevant” is more precisely defined in chapter two). Another convention, acknowledging the existence of a *programme counter* at a fixed locator, but hiding it, allows certain update schemes to be written as so-called *commands*. Any command exhibiting the pattern

$$PC[pc]p \ pc[OP \ args]qc \ \dots \ = [\dots] \Rightarrow PC[pc'] \ pc'[next] \ \dots .$$

may be written

$$OP \ args \ \dots \ = [\dots] \Rightarrow next \ \dots .$$

as, for example, in the update schemes in example 4 which may be part of the description of some zero-address machine. (This example anticipates some of the syntactic sugar introduced in chapter two.)

Example 4

$$\begin{aligned} \text{PUSH } x \ S[q] &\Longrightarrow S[p] \ p[x]q. \\ \text{ADD } S[q] \ [x \ y]q &\Longrightarrow S[p] \ p[x + y]q. \end{aligned}$$

These conventions are explained in more detail in chapter two. The command style of writing update schemes will, in chapter seven, be shown to be a special case of the archetype mechanism.

4 Overview

This thesis constitutes a definition of, and tutorial in Update Plans.

Chapters two and three formally define basic Update Plans, chapter two covering the syntax and syntactic sugar, and chapter three the semantics. Both of these chapters are illustrated by examples, as are chapters seven and nine in which, respectively, a macro-like mechanism and a degree of synchronous parallelism are added to the formalism.

A longer example, in which an abstract machine for an extended λ -calculus is defined is to be found in chapter four.

A typing mechanism for Update Plans is defined in chapter five. One application of typing is in detecting certain common types of memory use, and this is also covered in chapter five. These paradigms of memory use have well defined semantic properties. Some of these properties are presented in chapter six.

Basic Update Plans, as defined in chapters two and three, are well suited to the specification of abstract machines, but specifications of concrete machines can quickly become unwieldy due to the plethora of combinations of opcodes and addressing modes. The introduction of the archetype mechanism in chapter seven makes it easier to abstract away from this unwanted detail. Chapter eight, in which a register allocation algorithm is shown to preserve semantics, contains larger examples of applications of the archetype mechanism, and of the techniques developed in chapters five and six.

The archetype mechanism creates new possibilities, other than abstracting away from addressing modes, such as specifying asynchronous parallel processors. These possibilities are also covered in chapter seven. Another extension, allowing synchronous parallelism to be specified, is presented in chapter nine and illustrated in a specification of pipelining in the Berkeley RISC II machine.

It is intended that Update Plans be realisable as a full programming language, and as a first step in this direction some implementational aspects are reviewed in chapter ten.

Finally, chapter eleven contains the conclusions, and some suggestions for further research.

5 Notational Conventions

The following notational conventions will be observed throughout this thesis.

- Chapter numbers are always spelt out. All other reference numbers are expressed as numerals. This is, for example, ‘section 5 of chapter one’. Sections, figures, examples, etc. are all numbered within chapters. When referring to a section, figure, example, etc. not in the current chapter, the chapter reference will be given explicitly.
- In the running text single quotes are used to indicate strings, double quotes to indicate concepts — i.e. ‘string’ represents the sequence of symbols ‘s’, ‘t’, ‘r’, ‘i’, ‘n’ and ‘g’, and is an

example of a “string”.

- A typewriter font is used for update schemes. An italic font in an update scheme, e.g. *OP args* in the update schemes presented during the discussion of command style schemes on page 6, indicates a “meta-variable”.
- Concatenation of terms to form sequences is usually implicit in Update Plans. When it is necessary to make it explicit the concatenation symbol ‘++’ will be used.

This chapter defines the basic syntax of Update Plans, and introduces some syntactic sugar which makes update plans more readable. An informal description of the semantics of Update Plans was given in chapter one. A formal semantics can be found in chapter three.

1 Basic Syntax

The basic syntax of Update Plans is given by the following grammar.

$$\begin{aligned} \langle \text{script} \rangle &\rightarrow \langle \text{configuration} \rangle \cdot \langle \text{plan} \rangle \\ \langle \text{plan} \rangle &\rightarrow \langle \text{scheme} \rangle^* \\ \langle \text{configuration} \rangle &\rightarrow \langle \text{locator expression} \rangle^* \\ \langle \text{scheme} \rangle &\rightarrow \langle \text{configuration} \rangle \langle \text{guard} \rangle \langle \text{configuration} \rangle \cdot \\ \langle \text{guard} \rangle &\rightarrow \text{=}[\langle \text{term} \rangle] \Rightarrow \\ \langle \text{locator expression} \rangle &\rightarrow \langle \text{locator} \rangle [\langle \text{text} \rangle] \langle \text{locator} \rangle \\ \langle \text{locator} \rangle &\rightarrow \langle \text{term} \rangle \\ \langle \text{text} \rangle &\rightarrow \langle \text{term} \rangle \end{aligned}$$

A term is an expression built from constants, variables and operators, or a regular expression over the set of terms — e.g. 0^* represents a sequence of zeros. A regular expression over constants is called a *semi-constant*. If the regular expression contains no closures it is a *finite semi-constant*.

This grammar will be amended and expanded as extensions to Update Plans are introduced. A complete grammar for Update Plans can be found in appendix II.

Whitespace is used to separate the elements of a sequence in the above expressions — i.e. elements in sequences are separated by whitespace,

rather than commas or any other symbol. There must not be any whitespace between a locator and its brace ('[' or ']'), since this would lead to ambiguities. A tie character '~' may be used to override this, so that '...]~ L ...' is equivalent to '...]L ...'.

Rather than addressing the cells as if they were elements in an array, Update Plans address the *boundaries* between cells, so that what in an array notation would be specified as [2 : 4] = ['a', 'b', 'c'] is specified in an update scheme by 2['a' 'b' 'c']5. In this locator expression 2 is the *left locator* and 5 the *right locator* of the sequence of cells containing 'a', 'b' and 'c'. However, when no confusion can arise, a cell may be referred to by its left locator. That is, rather than saying "the cell with left locator X contains the value Y" one may say "X contains Y". In exceptional cases the right locator may be used in the same way, but then extra care must be taken to avoid confusion.

A singleton sequence is identified with its element. Everything in memory is considered to be a sequence, and singleton sequences are "desequenced" (to a locator, number or other basic type) as the context requires.

A variable is indicated by a lower case identifier. Constants are given by a value or, symbolically, by upper case identifiers. An identifier must start with a letter, and may comprise letters, digits, and the symbols '_' and "'". Identifiers may be subscripted. Thus 5, 's' and PC are all valid constants, x, new_pc and i' are all valid variable names and Num and TeX are neither. Constant locators play such a special rôle in Update Plans that a special, somewhat suggestive terminology is used to refer to them: a constant locator is called a *register*. By application of the convention that a cell may be referred to by its locator, a cell having a register as left or right locator may also be referred to as a register, if no confusion may arise as to which cell is being referred to.

2 Sweetening the Syntax

An update scheme conforming to the basic syntax can become unwieldy, as shown in example 1(a). While correct, the update schemes in example 1(a) lose legibility due to an excess of detail. Some extensions to the basic syntax will be introduced which provide abstraction capabilities and improve the readability of Update Plans. It is recommended that the reader first consider the simplified versions before spending too much time on the specification in example 1(a).

————— Example 1 (a) —————

The aim of this example is to construct a parser for a context free grammar. The grammar used in the example is

$$\langle S \rangle \rightarrow \mathbf{a} \mid \mathbf{a} \langle S \rangle$$

A script consisting of the following initial configuration

```

PARSE[START]PARSE' START[S]END END[STOP]END'
IP[I]IP' I[ω]EOF

```

where ω is some string, and the three update schemes

$$\begin{aligned} & \text{PARSE}[v]\text{PARSE}' v[S]w \\ & \Rightarrow [\text{TRUE}] \text{PARSE}[v']\text{PARSE}' v'['a']w. \end{aligned}$$

$$\begin{aligned} & \text{PARSE}[v]\text{PARSE}' v[S]w \\ & \Rightarrow [\text{TRUE}] \text{PARSE}[v']\text{PARSE}' v'['a']v'' v''[S]w. \end{aligned}$$

$$\begin{aligned} & \text{PARSE}[v]\text{PARSE}' v['a']w \text{ IP}[ip]\text{IP}' ip['a']ip' \\ & \Rightarrow [\text{TRUE}] \text{PARSE}[w]\text{PARSE}' \text{IP}[ip']\text{IP}'. \end{aligned}$$

will have a final development containing

$$\text{PARSE}[\text{END}]\text{PARSE}' \text{IP}[\text{EOF}]\text{IP}'$$

if and only if ω is in the language generated by S in the given grammar.

The extensions to the basic syntax fall into two categories. Firstly there is the syntactic sugar — notational changes which increase readability without introducing new capabilities. Secondly there are those notational conventions which also extend the expressive power of Update Plans.

3 Syntactic Sugar

There are four flavours of syntactic sugar, which will be illustrated by their application to the update scheme

$$\begin{aligned} & \text{Example 1 (b)} \\ & \text{PARSE}[v]\text{PARSE}' v[S]w \\ & \Rightarrow [\text{TRUE}] \text{PARSE}[v']\text{PARSE}' v'['a']v'' v''[S]w. \end{aligned}$$

- Superfluous locators may be omitted. A locator is superfluous if its removal does not lead to any confusion. Using terminology to be introduced in chapter five, the locator's removal may not alter the status with respect to grounding of any other expressions in the update scheme. In the following only right locators have been removed, but left locators may of course, under the same conditions, also be omitted.

$$\begin{aligned} & \text{Example 1 (c)} \\ & \text{PARSE}[v] v[S]w \\ & \Rightarrow [\text{TRUE}] \text{PARSE}[v'] v'['a']v'' v''[S]w. \end{aligned}$$

- Contiguous sequences may be concatenated. Two expressions $x[s]y$ and $y[t]z$ may then be written as $x[s]y[t]z$.

$$\begin{aligned} & \text{Example 1 (d)} \\ & \text{PARSE}[v] v[S]w \\ & \Rightarrow [\text{TRUE}] \text{PARSE}[v'] v'['a']v''[S]w. \end{aligned}$$

A further refinement is that $?[input]$ can be used for the standard input stream, when this is unambiguously defined.

Note that if the input pointer is not moved — the input is looked at, but not read — then this convention cannot be used.

There is a similar convention for output, using the symbol '!'. Any scheme with the pattern

$$\dots OP[o] \dots \implies \dots o[output]p OP[p] \dots .$$

may be written as

$$\dots \implies \dots !OP[output] \dots .$$

Again $![output]$ may be used for the standard output stream, if this is defined. The scheme in example 1(h) can now be written:

Example 1 (i)

$$\text{PARSE}[v] \ v['a']w \ ?['a'] \implies \text{PARSE}[w].$$

Programme counter. The update schemes obtained by applying all the conventions above satisfy the requirements, namely specifying a parser for the given grammar, but the introduction of new registers, such as `PARSE`, for each routine one wants to implement leads to illegibility and makes re-use of specifications in disparate contexts difficult. This can be simplified by introducing a command stream and a programme counter to administer it.

Example 1 (j)

The schemes above are summarised here for convenience.

$$\begin{aligned} \text{PARSE}[v] \ v[S]w &\implies \text{PARSE}[v'] \ v'['a']w. \\ \text{"} &\implies \text{PARSE}[v'] \ v'['a' S]w. \\ \text{PARSE}[v] \ v['a']w \ ?['a'] &\implies \text{PARSE}[w]. \end{aligned}$$

These can be transformed to a command stream style by replacing the constant locator `PARSE` by the constant “command” `PARSE` giving

$$\begin{aligned} \text{PC}[pc] \ pc[\text{PARSE } v]qc \ v[S]w \\ \implies \text{PC}[pc'] \ pc'[\text{PARSE } v']qc \ v'['a']w. \\ \text{PC}[pc] \ pc[\text{PARSE } v]qc \ v[S]w \\ \implies \text{PC}[pc'] \ pc'[\text{PARSE } v']qc \ v'['a' S]w. \\ \text{PC}[pc] \ pc[\text{PARSE } v]qc \ v['a']w \ ?['a'] \\ \implies \text{PC}[pc'] \ pc'[\text{PARSE } w]qc. \end{aligned}$$

The initial configuration becomes

$$\begin{aligned} \text{PC}[GO] \ GO[\text{PARSE } \text{START}] \\ \text{START}[S]\text{END} \ \text{END}[\text{STOP}] \ \text{IP}[I] \ I[\omega]\text{EOF} \end{aligned}$$

The string ω is in the language generated by `S` if there is a final development satisfying

At first sight introducing a programme counter does not seem to have increased the legibility of the update schemes, but the following notational convention, already introduced in chapter one, puts that right.

If an update scheme exhibits the pattern

$$\begin{aligned} & PC[pc] \quad pc[command]qc \quad \dots \\ & \quad \quad \quad \Rightarrow PC[pc'] \quad pc'[next]qc \quad \dots \end{aligned}$$

where *command* and *next* are some, possibly empty, sequences, and the locators PC, pc, qc and pc' can be omitted without confusion, it may be rewritten as

$$OP \text{ args } \dots \Rightarrow next \dots$$

An update scheme of this type is called a *command*. The first term of *command* will usually be some constant, known as the *opcode*. The sequences are known as *command sequences*. An update plan in which all update schemes are commands is called a *command driven* plan. The convention may also be applied individually to the left or right hand side, if the other side is not in command form. A configuration is in *command form* if it contains a non-empty command sequence, or if the contents of PC (the contents of the cell with left locator PC) are not specified. When only one side of an update scheme need be desugared the locator qc is not shared.

In most cases it will be necessary to apply this convention to the left and right hand sides simultaneously in order to satisfy the conditions under which locators may be omitted. The PARSE update scheme under consideration is a case in point. Rewriting only the left hand side gives

$$PARSE \ v \ v[S]w \Rightarrow PC[pc'] \quad pc'[PARSE \ v']qc \ v'['a' \ S]w.$$

which contains the undefined (or non-ground, see chapter five) locators pc' and qc. Similarly, rewriting only the right hand side gives

$$PC[pc] \quad pc[PARSE \ v]qc \ v[S]w \Rightarrow PARSE \ v' \ v'['a' \ S]w.$$

which, when desugared to

$$\begin{aligned} & PC[pc] \quad pc[PARSE \ v]qc \ v[S]w \\ & \quad \quad \quad \Rightarrow PC[pc'] \quad pc'[PARSE \ v']qc' \ v'['a' \ S]w. \end{aligned}$$

contains the non-ground locators pc' and qc'. Simultaneously rewriting both sides gives the update plan

Example 1 (k)

$$\begin{aligned} PARSE \ v \ v[S]w & \Rightarrow PARSE \ v' \ v'['a']w. \\ " & \Rightarrow PARSE \ v' \ v'['a' \ S]w. \end{aligned}$$

$$PARSE \ v \ v['a']w \ ?['a'] \Rightarrow PARSE \ w.$$

Detecting command style update schemes is (almost) a trivial syntactic exercise. If a non-empty command is present it must be desugared as described above. An empty command is assumed to be present if the contents of PC are not specified. Note that the contents of PC may also be specified when a non-empty command is present, making the contents of PC available in a command style update scheme. When both the left and right hand side are desugared the right locators of both commands must be identical. The following examples should make the desugaring mechanism clear.

Example 2

The update scheme

$$\text{JUMP } 1 \implies \text{PC}[1].$$

becomes

$$\text{PC}[\text{pc}] \text{ pc}[\text{JUMP } 1]\text{qc} \implies \text{PC}[1].$$

The left hand side is desugared, the right hand side isn't, since there is no non-empty command, and the contents of PC are specified. The update scheme:

$$\text{CALL } 1 \text{ PC}[\text{p}] \text{ SP}[\text{t}] \implies \text{PC}[1] \text{ SP}[\text{s}] \text{ s}[\text{p}]\text{t}.$$

becomes

$$\begin{aligned} \text{PC}[\text{pc}] \text{ pc}[\text{CALL } 1]\text{qc} \text{ PC}[\text{p}] \text{ SP}[\text{t}] \\ \implies \text{PC}[1] \text{ SP}[\text{s}] \text{ s}[\text{p}]\text{t}. \end{aligned}$$

Again the left hand side is desugared and the right hand side isn't. Note that consistency of the instantiation of the left hand side insures that p is equal to pc.

$$\implies \text{SILLY}.$$

becomes

$$\text{PC}[\text{qc}] \implies \text{PC}[\text{pc}] \text{ pc}[\text{SILLY}]\text{qc}.$$

Both sides are desugared. The left hand side contains the empty command.

Alternatives. An update plan may contain update schemes implementing some form of case analysis — a set of update schemes having part or all of their left hand sides in common, and distinguished from one another by mutually exclusive implicit or explicit guards. This may often entail repetitive occurrences of these guards, as in

$$\begin{aligned} lhs_1 & \quad = [g_1] \Rightarrow \quad rhs_1. \\ lhs_2 & \quad = [\neg g_1 \wedge g_2] \Rightarrow \quad rhs_2. \\ & \quad \vdots \\ lhs_n & \quad = [\neg g_1 \wedge \neg g_2 \wedge \dots \wedge g_n] \Rightarrow \quad rhs_n. \end{aligned}$$

Such schemes may be written as *alternatives*, in which only the first applicable update scheme, reading from top to bottom, will be applied. Alternatives are separated by semi-colons. The schemes above can be written

$$\begin{array}{l} lhs_1 \quad = [g_1] \Rightarrow \quad rhs_1; \\ lhs_2 \quad = [g_2] \Rightarrow \quad rhs_2; \\ \quad \quad \quad \vdots \\ lhs_n \quad = [g_n] \Rightarrow \quad rhs_n. \end{array}$$

Alternatives can be eliminated from an update plan by replacing the guards that have been omitted. Implicit guards will be detected by the typing mechanism (see chapter five).

5 Further Simplifications

As an aside, note that the update schemes in example 1(k) can be simplified further, at the cost of a simple change to the initial configuration. The first step is to eliminate the indirection in PARSE's argument, which gives the update script shown in example 1(l).

Example 1 (l)

Elimination of indirection gives the update plan

$$\begin{array}{l} \text{PARSE S} \quad \Rightarrow \quad \text{PARSE 'a'}. \\ \quad \quad \quad \text{"} \quad \quad \Rightarrow \quad \text{PARSE 'a' S}. \\ \\ \text{PARSE 'a' ?['a']} \Rightarrow \text{PARSE}. \end{array}$$

The end of the parse now has to be indicated in the initial configuration by modifying this to

$$\begin{array}{l} \text{PC[START] START[PARSE S STOP]} \\ \text{IP[I] I[\omega] EOF}. \end{array}$$

A final development must now satisfy

$$\text{PC[pc] pc[PARSE STOP] IP[EOF]}.$$

The PARSE command has now become superfluous since both S and 'a' are constants and can therefore both be used as opcodes.

Example 1 (m)

$$\begin{array}{l} \text{S} \quad \Rightarrow \quad \text{'a'}. \\ \quad \quad \quad \text{"} \quad \Rightarrow \quad \text{'a' S}. \\ \\ \text{'a' ?['a']} \Rightarrow \quad . \end{array}$$

The initial configuration is

$$\text{PC[START] START[S STOP] IP[I] I[\omega] EOF}.$$

A final development should satisfy

$$\text{PC[pc] pc[STOP] IP[EOF]}.$$

6 Summary

The syntax of Update Plans, with syntactic sugar as defined in this chapter, is given by the following context free grammar. This grammar makes use of the notational conventions of the specification formalism ASF+SDF [36, 37] in which the notation $\{S \text{ term}\}^+$ indicates a list of one or more S 's separated by the terminal term . If the $+$ is replaced by a $*$ the list may also be empty. The suffix *opt* indicates zero or one occurrences of its nonterminal.

$$\langle \text{script} \rangle \rightarrow \langle \text{configuration} \rangle \cdot \langle \text{plan} \rangle$$
$$\langle \text{plan} \rangle \rightarrow \langle \text{item} \rangle^*$$
$$\langle \text{item} \rangle \rightarrow \langle \text{alternatives} \rangle \cdot$$
$$\langle \text{alternatives} \rangle \rightarrow \{ \langle \text{scheme} \rangle ; \}^+ \cdot$$
$$\langle \text{scheme} \rangle \rightarrow \langle \text{configuration} \rangle \langle \text{guard} \rangle \langle \text{configuration} \rangle | \\ \langle \text{repeat} \rangle \langle \text{guard} \rangle \langle \text{configuration} \rangle$$
$$\langle \text{configuration} \rangle \rightarrow \langle \text{text} \rangle \langle \text{context} \rangle$$
$$\langle \text{text} \rangle \rightarrow \langle \text{term} \rangle^*$$
$$\langle \text{context} \rangle \rightarrow \langle \text{locator expression} \rangle^*$$
$$\langle \text{repeat} \rangle \rightarrow "$$
$$\langle \text{guard} \rangle \rightarrow = [\langle \text{term} \rangle] \Rightarrow | \Rightarrow$$
$$\langle \text{locator expression} \rangle \rightarrow \\ \langle \text{left-section} \rangle^+ \langle \text{locator} \rangle | \\ ? \langle \text{term} \rangle\text{-opt} [\langle \text{contents} \rangle] | \\ ! \langle \text{term} \rangle\text{-opt} [\langle \text{contents} \rangle]$$
$$\langle \text{left-section} \rangle \rightarrow \langle \text{locator} \rangle [\langle \text{contents} \rangle]$$
$$\langle \text{locator} \rangle \rightarrow \langle \text{term} \rangle\text{-opt}$$

The basis for Update Plans is a formal notation — locator expressions — for specifying machine configurations. These are then combined in update schemes, which specify possible machine transitions in which the machine is updated from a configuration specified by the left hand side of the update scheme to a configuration specified by the right hand side. Chapter two defined the syntax of Update Plans. In this chapter the semantics of Update Plans is formally defined.

This chapter consists of two sections. Section 1 discusses representations of computer memory and relates locator expressions to machine configurations. In section 2 this relation is extended to cover the full syntax of Update Plans.

1 Memory

Computer memory can be viewed in two complementary ways. In the first view, the machine on which update plans operate consists simply of a *memory* containing a finite number of stores, each consisting of a countable set of *cells* addressed by a completely ordered set of *locators*. Since these multiple stores can easily be mapped to a single store the machine will in the following be considered to consist of only one store.

In the second view, the memory is a function \mathcal{M} from locators to values. The domain of \mathcal{M} , or address space \mathcal{A} , is a countable set of locators. Since it is a countable set, there is a complete ordering $<_{\mathcal{A}}$ along with the relevant successor and predecessor functions $\text{succ}_{\mathcal{A}}$ and $\text{pred}_{\mathcal{A}}$. In update plans ‘+’ and ‘-’ are used in the usual way: $L + n$ to indicate $\text{succ}_{\mathcal{A}}^n L$ and $L - n$ for $\text{pred}_{\mathcal{A}}^n L$. The address space is two way infinite, that is $\text{succ}_{\mathcal{A}}$ and $\text{pred}_{\mathcal{A}}$ are defined for all elements of \mathcal{A} .

1.1 Locators

In the following the term ‘memory’ will be used to refer to the whole memory of the machine. This means that, in the formal view, \mathcal{M} is a total function, possibly returning some special value \perp for undefined values. A finite subset of the memory function will be referred to as a ‘configuration’.

A configuration is *satisfied* by any memory or configuration of which it is a subset, and only by such a configuration or memory. The domain of a configuration c — the subset of \mathcal{A} for which c is defined — is given by $\text{Dom } c$. Update schemes define relations between configurations. An application of an update scheme to a memory defines a relation between memories.

A locator expression defines a configuration. The locator expression $L[S]L'$ expresses the idea “the cells between L and L' contain the sequence S ”. Since the locator expression contains no information about the cells other than those between L and L' , only the values in the cells between L and L' are defined. More formally, there is a function I from locator expressions to configurations given by

$$I[l[s_0 \dots s_n]r] = \{(l + i, s_i) \mid 0 \leq i \leq n \wedge l + n = r\}.$$

Locator expressions can be combined to specify larger configurations. A set of locator expressions only specifies a configuration if it is *consistent*, i.e. there are no locator expressions in the set specifying conflicting contents for one and the same cell. For example $3[1\ 2]5\ 4[2\ 3]6$ is consistent, but $3[1\ 2]5\ 4[1\ 2]6$ is not, due to the conflict between $4[2]5$ and $4[1]5$.

Consistency can be defined formally as follows. Two configurations c and c' are consistent if $c\ x = c'\ x$ for all x in $(\text{Dom } c) \cap (\text{Dom } c')$. A set of locator expressions L is consistent if $I[l]$ is consistent with $I[l']$ for all l and l' in L .

1.2 Combining Configurations

Configurations can be combined in a manner akin to set union. The “update union” of two configurations must also be a configuration, i.e. a partial function. If the two configurations to be combined are not consistent then, by definition, there is an argument for which they give different values. One of the configurations is chosen to consistently provide the result in such cases. By convention this is the left argument of the “update union” operator.

Formally

$$c \uplus c' = c \cup (c' \upharpoonright (\text{Dom } c' \setminus \text{Dom } c)),$$

where $f \upharpoonright S$ is the restriction of the function f to the set S . This definition is equivalent to

$$\begin{aligned} (c \uplus c')\ l &= c\ l, \text{ if } l \in \text{Dom } c \\ &= c'\ l, \text{ if } l \notin \text{Dom } c \end{aligned}$$

Note that \uplus is equivalent to conventional set union if the configurations to which it is applied are mutually consistent.

The union operator defined above is used to extend the function $I[\cdot]$ to cover collections of locator expressions. This extension is noted $\mathcal{I}[\cdot]$, and is defined by

$$\mathcal{I}[e_1 \dots e_n] = I[e_1] \uplus I[e_2] \uplus \dots \uplus I[e_{n-1}] \uplus I[e_n]$$

The motivation for this definition can be found in section 2.

2 Update Schemes

Locator expressions are the basic building blocks of update schemes. The locator expressions presented above contain only constants, or instantiated values such as `3` and `'b'`. This is not the case in update schemes. These may contain variables, or uninstantiated values.

A locator expression containing variables is instantiated by substituting values for the variables to obtain a closed locator expression. A *substitution* is a mapping from variables to values. A value is a closed term. Given an expression e and a substitution σ , the instantiation of e by σ is noted e^σ .

An update scheme is constructed from two sets of locator expressions, forming the left and right hand sides, and a guard.

Example 1

Consider the following update scheme:

$$\text{PC}[x] \quad x[\text{BCC } o]y \quad \text{C}[0] \quad \Rightarrow \text{PC}[y+o].$$

This is a definition for the BCC instruction on the M6800 processor. Intuitively the definition states that if the instruction addressed by the PC is BCC with argument o and the condition code register C bit is 0 then the PC is reset to $y + o$ where y is the address immediately after the BCC o instruction.

More formally, this update scheme can be applied to the memory \mathcal{M} if the following conditions hold.

$$\mathcal{M} x = \text{BCC}, \text{ where } x = \mathcal{M} \text{ PC}$$

$$\mathcal{M} c = 0, \text{ where } c = \mathcal{M} \text{ C}$$

or equivalently $((\mathcal{M} \circ \mathcal{M}) \text{ PC} = \text{BCC}) \wedge (\mathcal{M} \text{ C} = 0)$. If these conditions hold then the result of applying the update scheme is a new memory \mathcal{M}' defined as follows.

$$\mathcal{M}' \text{ PC} = y + o$$

$$\text{where } y = x + 2$$

$$o = \mathcal{M}(x + 1)$$

$$x = \mathcal{M} \text{ PC}$$

$$\mathcal{M}' x = \mathcal{M} x, \text{ if } x \neq \text{PC}$$

An update scheme is *applicable* to a configuration if there is a substitution under which the guard is true, and the left hand side of the update scheme is contained in the configuration. The result of applying the update scheme is then this configuration, minimally changed so as to contain the instantiation of the right hand side, under the same substitution as that applied to the left hand side and the guard. The precise conditions under which such a substitution can be found will be discussed in chapter five, in which *grounding* is introduced, but, intuitively, a variable is ground if a closed expression can be derived for it. In fact a less restrictive requirement, *semi-grounding*, is defined in chapter five, but for this example the concept of grounding is sufficient.

The following example should clarify these concepts. Consider the update scheme for the BCC instruction of example 1, and the configuration $\text{PC}[1234] \ 1234[\text{BCC } 67]1236 \ \text{C}[0]$. A substitution, $\{(x, 1234), (y, 1236), (o, 67)\}$, can be found under which the left hand side is contained in, and in fact is identical to the given configuration, and the update scheme is therefore applicable to the configuration. The result is then the configuration $\text{PC}[1303] \ 1234[\text{BCC } 67]1236 \ \text{C}[0]$. Only the contents of PC have been changed, this being the minimal change necessary to make the configuration satisfy the right hand side of the update scheme under the substitution above. The update scheme is also well formed. All variables occurring in it are ground. An expression for o , for example, is $\mathcal{M}((\mathcal{M} \ \text{PC}) + 1)$.

More formally, an update scheme $(lhs, guard, rhs)$ is applicable to a configuration c if there is a substitution σ such that rhs^σ is consistent, $lhs^\sigma \subseteq c$ and $guard^\sigma$ holds. The result is the configuration c' given by

$$c' = \mathcal{I}[\![rhs^\sigma]\!] \uplus c.$$

In this configuration one can distinguish three types of cells. There are those cells specified in the configuration c , but not in the right hand side of the update scheme. The values in these cells remain unchanged. Secondly, there are the cells specified both in c and the right hand side of the scheme. The values in these cells are *updated* to the values specified in the right hand side, though they need not, of course, actually change. Finally, there may be some cells specified in the right hand side of the update scheme that do not occur in c . These cells are then added to the configuration.

The motivation for the definition of \uplus , and its application in the definition of $\mathcal{I}[\![\cdot]\!]$, should now be clear. The left argument of \uplus specifies an update of the right argument. Operationally, any cells in memory specified in the left configuration must have their value updated to the value specified in that configuration. The values in cells not specified in the left argument are left unchanged.

A substitution must now be found. The substitution must satisfy the guard and, when applied to the left hand side, it must give a configuration contained in the configuration to which the update scheme is being applied. More formally, a set of locator expressions L , a guard g and a configuration c together define a set of applicable substitutions as given by the function \mathcal{S} .

$$\mathcal{S} \ L \ g \ c = \{\sigma \mid (\mathcal{I}[\![L^\sigma]\!] \subseteq c) \wedge g^\sigma\}$$

Note that consistency of L under the substitution is not explicitly required. However, if c is consistent the instantiation of L must, implicitly, also be consistent.

The concept of applying update schemes can now be formalised. An update scheme $s = (lhs, guard, rhs)$ defines a relation between configurations, notation \Rightarrow_s , given by

$$c \Rightarrow_s c' \equiv \exists \sigma \in (\mathcal{S} \ lhs \ guard \ c) : rhs^\sigma \text{ is consistent} \wedge c' = \mathcal{I}[\![rhs^\sigma]\!] \uplus c.$$

When no ambiguity can arise the subscript s may be omitted. The relation \Rightarrow_s defines all configurations c' that can be derived from configuration c by application of update scheme s . There must be a substitution, delivered by the function \mathcal{S} , under which the left hand side of s is a subset of c , and under which the guard of s is true. The result of applying s is then c , minimally updated to be consistent with the instantiation, under the same substitution, of the right hand side of s .

One update scheme is typically used to specify one possible machine transition, for example, the effect of one specific instruction. A complete machine specification, such as that in chapter four, will contain many such update schemes. The transition relation defined above can easily be extended to update plans. Any update scheme, contained in the update plan, which is applicable to a given configuration may be applied to that configuration. An update plan P defines the following relation.

$$c \Rightarrow_P c' \equiv \exists s \in P : c \Rightarrow_s c'.$$

In \Rightarrow_P a nondeterministic choice is made of which update scheme to apply. Again, the subscript may be omitted when no confusion can arise. As usual \Rightarrow^* is the reflexive transitive closure of \Rightarrow .

A configuration c' is a *development* of configuration c under an update plan P if $c \Rightarrow_P^* c'$. A *final development* of a script consisting of initial configuration c and update plan P is any configuration c' such that $c \Rightarrow_P^* c'$ and $\neg \exists c'' : c' \Rightarrow_P c''$. To arrive at a final development one keeps on applying update schemes until one arrives at a configuration to which no update schemes in the plan are applicable — i.e. a final development is analogous to a normal form in a term rewrite system.

Chapter Four An Introductory Application and Example

In this chapter a slightly more complex example than those already presented is given. In this example an abstract machine, based on the λ MMachine [52], is specified in terms of an update plan. The λ MMachine is an abstract machine for the λ -calculus. It is a register based graph reduction machine using environments. The machine has been extended to handle constants, strict abstraction, some predefined functions and simple i/o. This extended λ -calculus is called the λ^+ -calculus. The λ^+ -calculus machine, or λ^+ MMachine, is provided with some syntactic sugar, providing global and local function declarations and infix operators. The result is a simple functional language called ‘FLIP’, for ‘Functional Language Implementation Prototype’. The acronym can also be seen as describing the four stages of compilation, as shown in figures 1 and 2:

- **F**unctions (or **FLIP** itself)
- **L**ambda calculus
- **I**ntermediate code
- **P**rogramme (or update **P**lan)

The λ MMachine was chosen as an example for several reasons.

- Functional languages have a simple syntax — a complex syntax would unnecessarily complicate the presentation.
- Functional languages present non-trivial implementation problems.
- Update Plans are intended for specifying those implementational details, primarily of a memory management nature, which would only obfuscate the specification given by Meijer and Patterson [52].

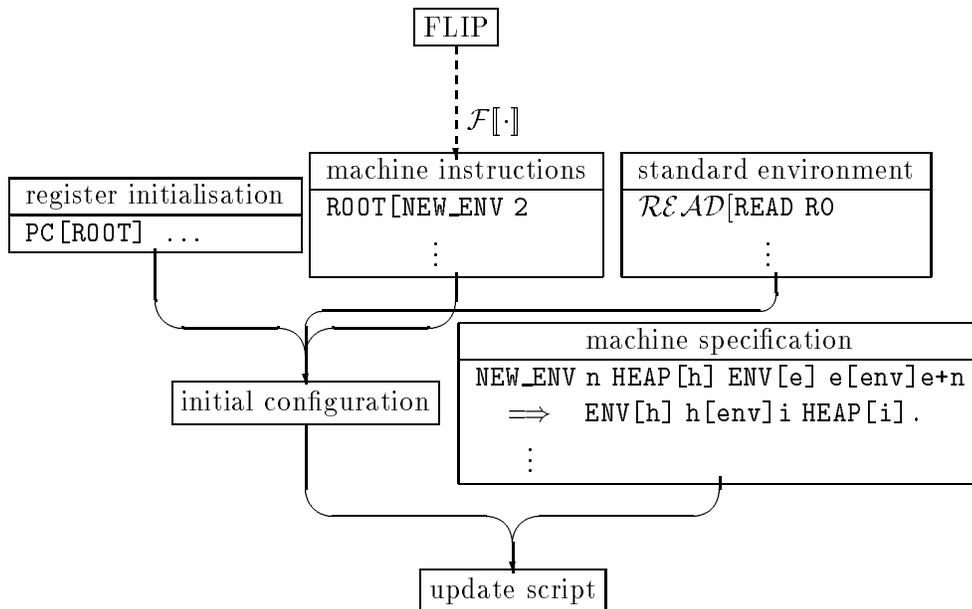


Figure 1: Compiling FLIP, Code Sequence to Update Script

The following section discusses compiling λ MMachine programmes to FLIP machine code, and gives an overview of the remainder of this chapter.

1 Compiling λ MMachine Programmes to FLIP code

A FLIP programme is compiled to a configuration — a specification of the contents of the programme store — for the FLIP machine. This is then combined with programme independent initialisations — register initialisations (given on page 28) and a standard environment — to form an initial configuration. This initial configuration, in combination with the update schemes specifying the instruction set of the FLIP machine (summarised in appendix III), forms an update script, execution of which will terminate in a configuration containing a representation of the value of the original FLIP programme. This is discussed in more detail below, and represented schematically in figure 1.

In figure 1 \mathcal{F} is the compilation function. \mathcal{F} is the composition of three auxiliary functions, \mathcal{D} , \mathcal{E} and \mathcal{T} , with $\mathcal{F}[f] = \mathcal{T}[\mathcal{E}[\mathcal{D}[f]] \sigma \tau] \text{ROOT}$, where σ and τ are lists of register names, τ being an infinite list of free registers, and σ a list of registers that have been pushed to a compile time stack. **ROOT** is the locator in **PC** in the initial configuration (see page 28). \mathcal{D} is not given explicitly. It is just a desugaring of FLIP syntax. \mathcal{E} translates a FLIP programme to intermediate code, and \mathcal{T} compiles this to FLIP machine code. A concrete example of this compilation path is given in figure 2. The example programme in figure 2 includes an application of strict abstraction, indicated by exclamation marks, and a predefined i/o function (**read**). These extensions to the λ -calculus will be covered in sections 4 and 6, respectively. The translation of the intermediate command *STORE* x to the code sequence **POP RO BIND 0 RO** is discussed

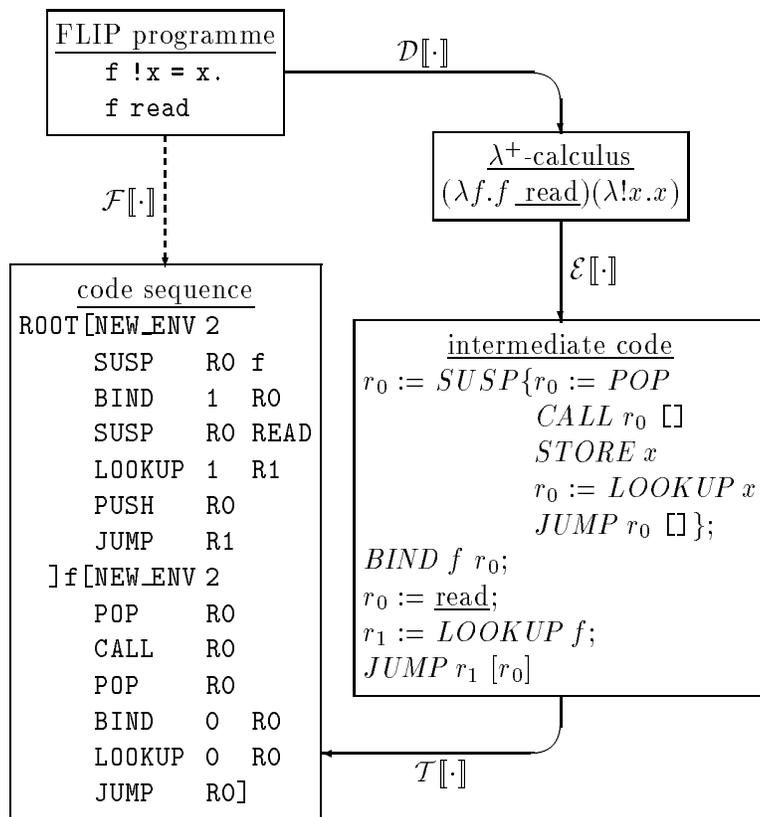


Figure 2: Compiling FLIP, FLIP to Code Sequence

in section 4.4.

\mathcal{F} is specified incrementally. Section 2 gives \mathcal{F} for the kernel of the language, section 3 defines the extension of \mathcal{F} for integers, section 4 that for strictness, section 5 for predefined functions, and section 6 completes the specification by defining \mathcal{F} for i/o.

Each section is structured in accordance with the compilation path, and will therefore contain, in the following order,

- a description of the (extension of the) syntax of the λ -calculus, which is the desugared version of (the extension to) FLIP.
- an informal description of the intermediate code produced by \mathcal{E}
- the function \mathcal{E} (for this extension)
- the translation function $\mathcal{T}[\cdot]$ from intermediate code to abstract machine instructions (again for this extension)
- update scheme specifications of the abstract machine commands which result from \mathcal{T} .

The update schemes will be numbered for ease of reference. Due to the incremental nature of the presentation it will sometimes be necessary to replace one update scheme with another. In this case the replacement will

not receive a new number, but will inherit the number of the scheme it replaces.

For any well formed FLIP programme f the code sequence $\mathcal{F}[[f]]$ forms part of the initial configuration of the desired update script. The remainder of the configuration, which is independent of the programme being compiled, consists of a standard environment, discussed in sections 5 and 6, and the following programme independent initialisations.

```
PC[ROOT]
HEAP[H] ENV[H] SP[S] S[END] END[STOP] TP[T] T[S].
```

The significance of the registers `HEAP`, `ENV`, `SP` and `TP` will be made clear during the specification of the machine. The specifications of the abstract machine commands, presented below and summarised in appendix III, form the update plan part of the script. The formation of the update script is shown schematically in figure 1. The script will be unambiguous and, applied if necessary to the input stream, its final development will be a representation of the value of the FLIP programme. Clearly, the value of the programme in figure 2, given input ‘2’, is 2. In the following specification a basic value, such as an integer, is represented by a *constructor* somewhere in the store. A constructor is a pair, consisting of the constant `CONST` and the basic value. A programme terminating with such a value will halt at a jump to a constructor containing the value (the reason for this construction can be found in the specification, and has to do with the strict abstraction mechanism). Given input ‘2’, the script containing the code in figure 2 will evaluate to a final development satisfying

```
PC[pc] pc[JUMP r] r[h] h[CONST 2].
```

2 The Basic λ -Calculus Machine

2.1 Syntax

The syntax of the λ -calculus is not very relevant to the example in hand, and assumed known.

2.2 Intermediate Code

The λ MMachine contains a heap, a stack, an unspecified number of registers and a static code space. A function application is encoded as a *suspension*, consisting of a pointer to a code sequence and a pointer to an environment, in which free variables are bound to their values.

The following intermediate code instructions are needed in order to implement this:

```
 $r := LOOKUP\ x$ 
    Look up the value of  $x$  (the suspension to which  $x$  is bound)
    in the current environment and assign this to register  $r$ .
 $r := SUSP\ \phi$ 
    Create a suspension consisting of a pointer to the code  $\phi$ 
    and a pointer to the current environment, and assign it to
 $r$ .
```

$r := POP$

Pop the topmost suspension from the stack and assign it to r .

$BIND\ x\ r$

Bind x in the current environment to the value of register r .

$JUMP\ r\ [r_1 \dots r_n]$

Ensure that the values of registers r_1 to r_n are on the stack, then jump to the code of the suspension bound to r while installing its environment.

2.3 $\mathcal{E}[\cdot]$, Translation to Intermediate Code

The translation scheme presented below is adapted from work done by Meijer and Patterson [52], as are the descriptions of the intermediate code.

Arguments of functions are conceptually pushed onto the stack, to be popped when needed. In order to improve efficiency these arguments are, whenever possible, stored in registers rather than being transferred via the stack. In order to do this a compile time “stack” of registers, emulating a run time stack, is maintained. Only when a function is called must the contents of the registers on the compile time stack be pushed onto the run time stack.

The central translation scheme is that for expressions: $\mathcal{E}[e] \sigma \tau$. The translation schemes for definitions, expressions and suspensions are:

$$\begin{aligned} \mathcal{E}[e_1\ e_2] (r \ ++\ \sigma) \ \tau &= r := \mathcal{C}[e_2]; \\ &\quad \mathcal{E}[e_1] \ \sigma (r \ ++\ \tau) \\ \mathcal{E}[\lambda x. e] \ \sigma (r \ ++\ \tau) &= BIND\ x\ r; \\ &\quad \mathcal{E}[e] (r \ ++\ \sigma) \ \tau \\ \mathcal{E}[\lambda x. e] (r \ ++\ \sigma) \ \square &= r := POP; \\ &\quad BIND\ x\ r; \\ &\quad \mathcal{E}[e] (r \ ++\ \sigma) \ \square \\ \mathcal{E}[e] (r \ ++\ \sigma) \ \tau &= \mathcal{C}[e] : \\ &\quad JUMP\ r\ \tau \end{aligned}$$

where the auxiliary function \mathcal{C} is used to produce intermediate code for the arguments of an application. \mathcal{C} is:

$$\begin{aligned} \mathcal{C}[x] &= LOOKUP\ x \\ \mathcal{C}[e] &= SUSP\ \mathcal{E}[e] \ \Sigma \ \square \\ &\text{where } \Sigma \text{ is the list of all available registers.} \end{aligned}$$

See figure 2 for an example of an application of these and other translation schemes.

2.4 $\mathcal{T}[\cdot]$, Translation to Machine Configurations

The translation to machine configurations is not difficult. The function \mathcal{U} translates intermediate code to initial configurations. It returns a pair. The first element of the pair is a translation of the instructions occurring in the argument of \mathcal{U} ; the second is a translation of the suspensions. The function \mathcal{T} ensures that these are concatenated into a correct configuration.

The prototype presented here implements environments in a very simple minded way — each suspension in which new bindings are made creates its own copy of the universal environment, containing fields for all variables in the script. A more sophisticated implementation would probably use a combination of de Bruijn numbering [17] and some dereferencing mechanism. \mathcal{T} is:

$$\mathcal{T}[\mathit{code}] \mathit{loc} = \mathit{loc}[\mathbf{NEW_ENV} \ \Sigma \ \mathit{instrs}] \ \mathit{susps}$$

where Σ is the size of the universal environment, instrs is a code sequence for the function encoded in code , and susps is a set of locator expressions containing the code sequences of functions called in code — i.e. the code sequences of suspensions occurring in code . These two translations, instrs and susps are derived as follows. The function \mathcal{U} , defined below, when applied to an intermediate code command, returns a pair consisting of a FLIP machine code sequence implementing this intermediate code command and, for a $\mathit{SUSP} \ \mathit{susp} \ \mathit{r}$ command, a set of locator expressions implementing the suspension susp . When mapped across the intermediate code sequence code the function \mathcal{U} will therefore return a list of pairs of FLIP machine code sequences and lists of locator expressions. The pair is brought outside the list — i.e. the list of pairs of lists becomes a pair of lists of lists, and these lists of lists are flattened to give a pair of lists. The first element of the pair is then instrs and the second susps . A full definition of \mathcal{T} is:

$$\begin{aligned} \mathcal{T}[\mathit{code}] \mathit{loc} &= \mathit{loc}[\mathbf{NEW_ENV} \ \Sigma \ \mathit{instrs}] \ \mathit{susps} \\ \text{where } (\mathit{instrs}, \mathit{susps}) &= (\mathit{flatten} \ \mathit{instrs}', \mathit{flatten} \ \mathit{susps}') \\ (\mathit{instrs}', \mathit{susps}') &= \mathit{lift} \ (\mathit{map} \ \mathcal{U} \ \mathit{code}) \\ \mathit{lift} \ \mathit{prs} &= (\mathit{map} \ \mathit{fst} \ \mathit{prs}, \mathit{map} \ \mathit{snd} \ \mathit{prs}) \\ \mathit{fst} \ (a, b) &= a \\ \mathit{snd} \ (a, b) &= b \end{aligned}$$

\mathcal{U} is defined for the intermediate code instructions in the translation scheme for \mathcal{E} above as follows:

$$\begin{aligned} \mathcal{U}[\mathit{r} := \mathit{LOOKUP} \ x] &= (\mathbf{LOOKUP} \ \sigma(x) \ \mathbf{r}, \ \epsilon) \\ &\text{where } \sigma \text{ is the indexing function for identifiers in the} \\ &\text{environment.} \\ \mathcal{U}[\mathit{r} := \mathit{SUSP} \ \{\mathit{susp}\}] &= (\mathbf{SUSP} \ \mathbf{r} \ l, \quad \mathcal{T}[\mathit{susp}] \ l) \\ &\text{where } l \text{ is some unique locator} \\ \mathcal{U}[\mathit{r} := \mathit{POP}] &= (\mathbf{POP} \ \mathbf{r}, \quad \epsilon) \\ \mathcal{U}[\mathit{BIND} \ x \ \mathit{r}] &= (\mathbf{BIND} \ \sigma(x) \ \mathbf{r}, \quad \epsilon) \\ \mathcal{U}[\mathit{JUMP} \ \mathit{r} \ [\mathit{r}_1.. \mathit{r}_n]] &= (\mathbf{PUSH} \ \mathbf{r}_n \\ &\quad \vdots \\ &\quad \mathbf{PUSH} \ \mathbf{r}_1 \\ &\quad \mathbf{JUMP} \ \mathbf{r}, \quad \epsilon) \end{aligned}$$

Again, an example of an application of $\mathcal{T}[\cdot]$ can be found in figure 2.

2.5 Machine Instructions

The environmental commands are:

[2.1] NEW_ENV n HEAP[h] ENV[e] e[env]e+n

\implies ENV[h] h[env]i HEAP[i].

[2.2] BIND i r r[h] ENV[e] \implies e+i[h].

[2.3] LOOKUP i r ENV[e] e+i[h] \implies r[h].

The stack commands are straightforward.

[2.4] PUSH r r[h] SP[t] \implies SP[s] s[h]t.

[2.5] POP r SP[s] s[h]t \implies r[h] SP[t].

The last two commands are SUSP, which creates a suspension on the heap, and JUMP which jumps to such a suspension.

[2.6] SUSP r c ENV[e] HEAP[h]

\implies r[h] h[c e]i HEAP[i].

[2.7] JUMP r r[h] h[c e]

\implies PC[c] ENV[e].

3 Basic Values

Only values of the type *integer* will be considered here. The construction used is such that it is fairly simple to extend it to include other types.

3.1 Syntax

Integers are represented by their standard decimal notation.

3.2 Intermediate Code

There is obviously no need to look up the value of a constant in the environment, so a new intermediate code instruction is needed which simply assigns the value of the constant to a register.

$r := c$

Assign the constant c to register r .

3.3 $\mathcal{E}[\cdot]$

A constant is its own translation.

$\mathcal{C}[c] = c$

3.4 $\mathcal{T}[\cdot]$

Constants can also be represented by suspensions. They will be tagged to distinguish them from ordinary suspensions. A new command is therefore introduced, analogous to the SUSP command, to create tagged suspensions containing values.

$\mathcal{U}[r := c] = (\text{ASSIGN } r \ \mu(c)$
 $\text{CONST } r, \quad \epsilon)$

where μ is a compile time function returning the internal representation of its argument.

3.5 Machine Instructions

The assignment command is straightforward.

$$[3.1] \text{ ASSIGN } r \ n \ \Longrightarrow \ r[n].$$

The suspensions created by the `CONST` command are tagged to indicate that they are constructor suspensions.

$$[3.2] \text{ CONST } r \ r[n] \ \text{HEAP}[h] \\ \Longrightarrow \ h[\text{CONST } n]i \ r[h] \ \text{HEAP}[i].$$

Standard suspensions will now also have to be tagged to distinguish them from constructor suspensions.

$$[2.6] \text{ SUSP } r \ c \ \text{ENV}[e] \ \text{HEAP}[h] \\ \Longrightarrow \ r[h] \ h[\text{SUSP } c \ e]i \ \text{HEAP}[i].$$

The `JUMP` command must now distinguish between the two types of suspension. In the case of a constructor there is no evaluation to be done and execution should terminate. Providing an update scheme for a `JUMP` to a standard suspension, which must be evaluated as before, but not for a `JUMP` to a constructor suspension ensures that execution will halt when the machine encounters a constant, which can only happen when a (well-formed) λ^+ -expression has been fully evaluated.

$$[2.7a] \text{ JUMP } r \ r[h] \ h[\text{SUSP } c \ e] \ \Longrightarrow \ \text{PC}[c] \ \text{ENV}[e].$$

(An update scheme 2.7b will be defined later.)

4 Strict Abstraction

4.1 Syntax

The strict annotation symbol in the λ^+ -calculus is an exclamation mark. It should not be confused with the ‘!’ output convention, introduced in chapter two.

4.2 Intermediate Code

Two new instructions are needed in the intermediate code.

CALL $r \ [r_1 \dots r_n]$
Call the suspension bound to r . Save the current evaluation so that the callee can return upon terminating. Save the values of registers $[r_1 \dots r_n]$ on the stack.

STORE x
Bind x in the current environment to the value returned by the preceding *CALL*.

4.3 $\mathcal{E}[\cdot]$

A strictly abstracted expression must be evaluated before its value is bound to its variable. Since a function is about to be called the contents of the compile time stack registers must be pushed onto the run time stack. The removal from the stack of the values pushed by *CALL* is ensured by

continuing the translation of the expression with an empty compile time stack.

$$\begin{aligned} \mathcal{E}[\lambda!x.e] \sigma (r \ ++ \ \tau) &= \text{CALL } r \ \tau; \\ &\text{STORE } x; \\ &\mathcal{E}[e] \Sigma \ \square \\ \mathcal{E}[\lambda!x.e] \sigma \ \square &= r \ := \ \text{POP}; \\ &\text{CALL } r \ \square; \\ &\text{STORE } x; \\ &\mathcal{E}[e] \Sigma \ \square \end{aligned}$$

where, again, Σ is the list of all available registers.

4.4 $\mathcal{T}[\cdot]$

The machine configuration for these new commands is simple.

$$\begin{aligned} \mathcal{U}[\text{CALL } r \ [r_n \ \dots \ r_1]] &= (\text{PUSH } r_n \\ &\quad \vdots \\ &\quad \text{PUSH } r_1 \\ &\quad \text{CALL } r, \quad \epsilon) \end{aligned}$$

$$\begin{aligned} \mathcal{U}[\text{STORE } x] &= (\text{POP } \text{RO} \\ &\quad \text{BIND } \sigma(x) \ \text{RO}, \quad \epsilon) \end{aligned}$$

where σ is, again, the indexing function for identifiers in the environment.

A **STORE** command could have been introduced as the translation of **STORE**, but at this level it is observable that results are always returned via the stack. The use of **RO** to hold the value being fetched is safe, since **STORE** is always the first command after a **CALL**, which has just cleared all registers by pushing them to the run time stack. Execution must no longer terminate when a constant is encountered, but rather the machine must return to the caller, if any, and place a pointer to the constructor suspension on the stack for the use of the caller. This will be dealt with in the next section.

4.5 Machine Instructions

Strict abstraction will be presented in two steps. Before general strict abstraction is considered, the evaluation method for first order functions will be presented. A function returning a basic value will terminate at a **JUMP** to a constructor expression. Execution must then return to the caller. The caller must therefore have created a suspension containing the continuation address and the current environment. The address of this suspension is passed, like any other argument, via the stack.

$$\begin{aligned} [4.1a] \text{PC}[p] \ p[\text{CALL } r]q \ r[h] \ h[\text{SUSP } c \ e] \\ \quad \text{ENV}[f] \ \text{HEAP}[i] \ \text{SP}[t] \\ \implies \text{PC}[c] \ \text{ENV}[e] \ i[\text{SUSP } q \ f]j \\ \quad \text{HEAP}[j] \ s[i]t \ \text{SP}[s]. \end{aligned}$$

This ensures that the requisite information is available for the **JUMP** command.

$$\begin{aligned}
[2.7b] \text{ JUMP } r \ r[h] \ h[\text{CONST}] \ \text{SP}[s] \ s[i] \ i[\text{SUSP } c \ e] \\
\implies \text{PC}[c] \ \text{ENV}[e] \ s[h].
\end{aligned}$$

(Note that this scheme does not replace scheme 2.7a, but is an alternative to it.)

The key to the specification of the `CALL` command is the locator expression $i[\text{SUSP } q \ f]j$ on the right hand side of scheme 4.1a. Here `CALL` has encountered a suspension (locator expression $h[\text{SUSP } c \ e]$ on the left hand side). This suspension must now be evaluated by moving c , the pointer to its code, to the programme counter, and moving its environment pointer e to the current environment register. Execution of the current function must continue upon completion of the call. A new suspension representing the current state of evaluation of the function is therefore created — the locator expression $i[\text{SUSP } q \ f]j$ — and its address is pushed onto the stack.

If `CALL` encounters a constructor suspension this can be returned immediately.

$$[4.1b] \text{ CALL } r \ r[h] \ h[\text{CONST}] \ \text{SP}[t] \implies \text{SP}[s] \ s[h]t.$$

Clearly this mechanism will only work if functions always return basic values — i.e. functions are always fully evaluated.

The second step in the presentation of strict abstraction is allowing partial evaluation of functions — making higher order functions possible. For example, in the FLIP programme

```

add x y = x + y.
f !g x = g x.

```

```
f (add 2) 3.
```

evaluation of `add 2`, forced by `f` being strict in its first argument, will yield a suspension representing the function $(\lambda x.add \ 2 \ x)$. Care must be taken that the function called does not pop more arguments from the stack than it was given. A call must simulate the creation of a new stack, with a new stack bottom. To this end a `CALL` notes the current top of the stack on a tripwire stack, addressed via the register `TP`. The value at the top of the tripwire stack is therefore the address of the bottom of the “new” stack.

$$\begin{aligned}
[4.1a] \text{ PC}[p] \ p[\text{CALL } r]q \ r[h] \ h[\text{SUSP } c \ e] \\
\text{ENV}[f] \ \text{HEAP}[i] \ \text{SP}[t] \ \text{TP}[v] \\
\implies \text{PC}[c] \ \text{ENV}[e] \ i[\text{SUSP } q \ f]j \\
\text{HEAP}[j] \ s[i]t \ \text{SP}[s] \ \text{TP}[u] \ u[s]v.
\end{aligned}$$

`POP` must now check for stack underflow against the value stored at the top of the tripwire stack. If the stack has reached the tripwire it is time to halt evaluation of the current suspension and return to the caller. A new suspension must be created containing the partially evaluated expression, and this must be returned as the result.

$$\begin{aligned}
[2.5a] \text{ POP } r \ \text{SP}[s] \ \text{TP}[u] \ u[w] \ s[h]t \\
=_{[s < w]} r[h] \ \text{SP}[t].
\end{aligned}$$

$$\begin{aligned}
[2.5b] \text{ POP } & \text{PC}[p] \text{ SP}[s] \text{ TP}[u] \text{ u}[w] \text{ v}[s] \text{ h}[t] \\
& \text{ h[SUSP c e] ENV}[f] \text{ HEAP}[i] \\
=& [s=w] \Rightarrow \text{PC}[c] \text{ ENV}[e] \text{ TP}[v] \\
& \text{ i[SUSP p f]j HEAP}[j] \text{ s}[i].
\end{aligned}$$

Similarly, a JUMP encountering a constructor suspension must pop the tripwire stack.

$$\begin{aligned}
[2.7b] \text{ JUMP } & \text{r r}[h] \text{ h[CONST] TP}[u] \\
& \text{ u}[s] \text{ v}[s] \text{ i}[t] \text{ i[SUSP c e]} \\
\Rightarrow & \text{PC}[c] \text{ ENV}[e] \text{ TP}[v] \text{ s}[h].
\end{aligned}$$

The update schemes are now probably reaching the limit of their readability. Chapter seven defines a macro-like mechanism which can abstract away from detail, and greatly increase the readability of update schemes.

5 Predefined Functions

In the discussion of basic values in section 3 values of the type *integer* were used as a typical example. Here the addition function plays the same rôle. Other predefined functions can be treated analogously.

5.1 Syntax

At the surface language level infix expressions such as $\mathbf{x+y}$ are allowed. However these are desugared to expressions such as $(\lambda!x!y.+ x y) x y$ at the core language level.

5.2 Intermediate Code

Since the functions are predefined there is no need to generate code for a suspension — this can be taken, for example, from the standard environment.

$$r := f$$

Create a suspension for the predefined function f and assign it to register r .

Unsurprisingly, this is similar to the intermediate code generated for basic values. A predefined function can, after all, be considered to be a basic value.

5.3 $\mathcal{E}[\cdot]$

The translation of a predefined function is simply that function, though it may help to think of it as the address of a suspension containing the code for that function.

$$\mathcal{C}[[f]] = f$$

5.4 $\mathcal{T}[\cdot]$

For any predefined function occurring in a FLIP programme the relevant predefined suspensions must be included in the machine configuration.

In a final implementation all predefined suspensions would probably be included in a standard environment.

$\mathcal{U}[\![r := f]\!] = (\text{SUSP } r \ \pi(f), \epsilon)$
 where π is a compile time function returning the address,
 in the standard environment, of the code for function f .

The code for $+$ could be

```

ADD [POP   RO
      POP   R1
      ADD   R1   RO
      CONST RO
      JUMP  RO]

```

where *ADD* is an internal locator constant. The *JUMP* may look a little strange here, but since the *CONST* will have created a constructor suspension in *RO*, the *JUMP* will simply return to the caller. A single instruction could have been constructed and specified combining the effects of *CONST RO* and *JUMP RO*, but the choice was made to keep the number of instructions low.

5.5 Machine Instructions

The *ADD* command must now access the values in the constructor suspensions addressed by *RO* and *R1*, and perform the addition. As stated above, the *CONST* command will then create a new suspension to contain the result and the *JUMP* will return this to the caller.

[5.1] *ADD* *r1* *r0* *r1*[*h*] *h*[*CONST* *x*] *r0*[*i*] *i*[*CONST* *y*]
 $\implies r0[x+y]$.

6 I/O

6.1 Syntax

The only i/o commands are *read* and *write*. The surface language *write* is compiled to $(\lambda!x.\underline{\text{write}}\ x)$ at the core language level. For the sake of simplicity it will be assumed that only integers will be the subject of i/o operations. Note that *write* is strict in its argument. Its only effect is to write the value of its argument to the standard output stream, and is otherwise semantically equivalent to $(\lambda!x.x)$.

6.2 Intermediate Code, $\mathcal{E}[\cdot]$

The i/o commands *read* and *write* are in fact predefined functions, and are treated as such.

6.4 $\mathcal{T}[\cdot]$

Since *read* has arity 0 its predefined code is

```

READ [READ   RO
      CONST  RO
      JUMP   RO]

```

Similarly, *write* has arity 1 and predefined code

```

WRITE [POP   RO
      WRITE  RO
      JUMP   RO]

```

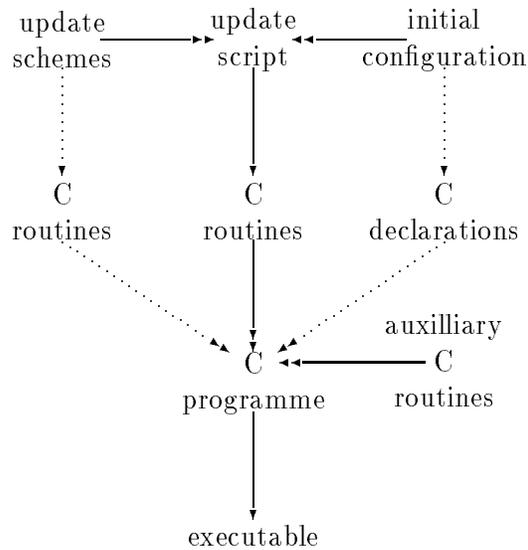


Figure 3: Compilation scheme for FLIP programmes

6.5 Machine Instructions

The constant locators ? and ! appearing in these schemes are the input and output buffers respectively. It is assumed that the i/o channels perform the necessary conversions.

[6.1] READ r ?[n] \implies $r[n]$.

[6.2] WRITE r $r[h]$ $h[\text{CONST } n]$ \implies ![n] .

7 The Implementation

The language and machine presented above have been implemented. A FLIP to abstract machine code compiler has been written. The machine configurations delivered by this compiler are compiled to C by a prototype Update Plan compiler. A complete Update Plan compiler would be capable of compiling update scripts, thus completing the compilation process sketched in figures 1 and 2. However, no such compiler is, as yet, available so the specifications of the commands given here were compiled by hand. This compilation was as close an imitation of machine compilation as possible.

The final stage of the FLIP compilation, update script to executable, is shown in figure 3. Single headed arrows indicate translation steps, double headed arrows simple concatenation. The solid lines in figure 3 indicate the compilation path that would be taken if there were a complete Update Plan compiler, the dotted lines the path actually taken to produce a prototype of the machine presented above. The resulting implementation is sufficiently efficient for realistic prototyping.

The efficiency of the code produced was tested by applying the FLIP compiler to a simple `nfib` programme. The `nfib` function is given by

`nfib` n = 1, if $n < 2$

$$= 1 + (\text{nfib } (n - 1)) + (\text{nfib } (n - 2)), \text{ otherwise}$$

and has the useful characteristic that `nfib` n is the number of applications of `nfib` necessary to compute `nfib` n .

The `nfib` number of a programme is the value of `nfib` n , divided by the number of seconds required to compute it. The higher the `nfib` number the better. Declarative language implementations have `nfib` numbers typically ranging from one of the order 1 [41] to 1,000,000 [11]. This last value is competitive with imperative languages. The gnu C compiler version 1.36 on a SUN 3/60 running OS 4.1 has an `nfib` number of 250,000.

The `nfib` number for this prototype of FLIP was of the same order as that for Miranda¹ (release 2). On a SUN 3/60, running under OS 4.1, Miranda had an `nfib` number of 1,200 and the FLIP implementation, compiled using the gnu C compiler version 1.36 with the optimisation flag set, 1,100. Simple optimisation of the intermediate code (eliminating unnecessary copying, for example) increased this to 1,400, and optimisation at the update plan level gave an `nfib` number of 2,200.

Though this may not appear impressive, it should not be forgotten that the `nfib` number for FLIP is for an implementation produced by a general prototyping system, while the `nfib` numbers given above are for dedicated implementations of specific declarative languages. The important point is that the `nfib` number is competitive with that of the Miranda interpreter, which experience has shown to be efficient enough for realistic programming. An `nfib` programme written directly in update schemes had an `nfib` number comparable to that of a programme written in C.

¹Miranda is a trademark of Research Software Ltd.

Two important aspects of Update Plans are their *implementability* and their *analysis*. Implementational aspects such as backtracking are dealt with in chapter ten, here “implementability” is intended to cover the more fundamental aspect of whether it is possible to implement a given update plan at all, in particular whether or not there is a manageable degree of ambiguity. This is related to the types of objects present in the update plan. Typing is also relevant to the analysis of memory use in update plans. Simple analysis generalises the requirements for certain standard types of memory structure such as stacks, heaps, input and output streams and static stores, thus providing a simple taxonomy of these structures.

Typing is introduced in section 1 and presented in detail in section 2. Section 3 discusses the relation between typing and ambiguity. The rôle of typing in detecting memory structures is covered in section 4, which defines some common memory structures. In section 5 *programme stores* are defined. The semantic properties of these, and other memory use paradigms, are discussed in chapter six.

1 Introduction

The basic type in Update Plans is the locator. In fact all objects appearing in an update plan are considered to be locators. One immediate consequence of this is that every type has a countable carrier set. In fact each type is the carrier of a unique countably infinite set. This does not exclude the use of update plans for specifications involving even simpler types, since any countable type can be realised by means of an injective function from its carrier to some set of locators. A disciplined use of this facility makes it possible to use Update Plans for very low level specifications, for example at bit sequence level by using reserved registers for conversion from bit sequences to locators.

The set of locators can also be partitioned into disjunct sets, or stores, as discussed in chapter three. Membership of one of these stores can also be considered as a form of typing, so the basic type *locator* is

actually a finite set of types. In this view any new type may, if its carrier set is countable, be defined as a new store having the elements of that carrier set as locators, and in this chapter ‘type’ and ‘store’ may be used interchangeably, as the context requires.

2 Typing

New types may be defined by combining existing types using any of the standard operators of regular expressions. Such a declaration is said to define a *type alias*. While some typing information may be derived automatically [45] it is the update plan writer’s responsibility to ensure that the plan is consistently typed. The writer must indicate which stores are present in the machine being specified. Store names are lower case words with an initial upper case letter, and must, therefore, contain at least two letters, in order to ensure that store names can be distinguished from constants. Stores are declared by listing them between braces:

```
{Stack, Heap, Int, Bool}
```

Each store name is said to be a *type primitive*. Type aliases are declared similarly:

```
{Programme = Routine  $\cup$  Coroutine}
```

Type aliases may not be recursive, either directly or indirectly. This ensures that any type alias can be expressed as a regular expression containing only type primitives. Every object (constant, variable, or expression) appearing in an update plan must be typed. For some objects this may be done implicitly. It is, for example, assumed that a number is of type `Int`, unless indicated otherwise. Each symbolic constant is considered to have its own unique type, again unless indicated otherwise. Other objects — expressions for which no type can be determined automatically — must have their type indicated. This can be done in two ways:

- by means of a global declaration

```
h :: Heap.
```

A global type declaration is valid throughout an update plan, unless overridden by a cast;

- or by *casting* a term within an update scheme,

```
... (h :: Heap)[...]...
```

such a cast determining only the type of the term to which it is applied. Any variables in the term share their value, coerced if necessary, with other occurrences of the variable in the update scheme. A cast is only properly defined if all necessary coercion mappings are defined.

A term given by a regular expression has the type defined by the same regular expression with all its constituent terms replaced by their types. The formal syntax of type declarations is obtained by modifying the grammar in chapter two as follows:

- Add the production rules:

$$\langle \text{item} \rangle \rightarrow \langle \text{store declaration} \rangle . \mid \langle \text{type declaration} \rangle .$$

- Add also the production rules:

$$\langle \text{term} \rangle \rightarrow (\langle \text{term} \rangle :: \langle \text{store structure} \rangle)$$

$$\langle \text{type declaration} \rangle \rightarrow \{ \langle \text{term} \rangle , \}^+ :: \langle \text{store structure} \rangle$$

- and the production rules:

$$\langle \text{store declaration} \rangle \rightarrow \{ \{ \langle \text{store} \rangle , \}^* \}$$

$$\langle \text{store} \rangle \rightarrow \langle \text{store name} \rangle = \langle \text{store structure} \rangle \mid \langle \text{store name} \rangle$$

A *store name* is a lower case word with a leading upper case letter, and a *store structure* is a regular expression over the set of store names.

Store declarations and the store structure parts of casts can also be seen as defining a context free grammar, known as the *type grammar*, having type names as its terminal symbols, and type aliases as its nonterminals. Some additional nonterminals may need to be introduced for implicit type aliases in casts. If a “pure” context free grammar is required yet more nonterminals may be needed to replace closure operators. Defining a context free grammar when type aliases can be defined in terms of regular expressions may seem an overkill, but the type grammar is useful as the basis for the *archetype grammar*, introduced in chapter seven.

Type declarations and casts associate a symbol from the type grammar to each variable appearing in the update plan.

3 Ground Expressions

An absolute requirement for implementability of Update Plans is that *bounded nondeterminism* be avoided. An update scheme exhibits unbounded nondeterminism if it may instantiate to infinitely many update rules, as in example 1 on page 42. Quite apart from the practical difficulties of providing any realistic implementation of Update Plans with unbounded nondeterminism, there are also serious theoretic reasons for avoiding it, since the presence of unbounded nondeterminism complicates any formal semantics [3, 4, 65].

Consider the update scheme

$$\Longrightarrow \mathbf{a}[0].$$

This update scheme will nondeterministically set the contents of some cell in the store to zero. If included in an update plan it will ensure that the store eventually (at time = ∞) contains nothing but zeros. Clearly this degree of nondeterminism cannot be permitted.

Unbounded nondeterminism is avoided by requiring every expression in an update scheme to be *ground* or *semi-ground*. Intuitively a ground expression is one for which a variable-free expression can be derived, possibly by instantiation of variables with respect to the current configuration; a semi-ground expression one for which a finite number of such expressions can be derived. A *statically semi-ground* term is one for which a variable free expression can be derived without reference to a configuration. A formal definition is given below.

Given typing

$$\begin{aligned} \mathbf{x} &:: \text{Some_type} . \\ \mathbf{xs}, \mathbf{ys}, \mathbf{zs} &:: (\text{Some_type})^* . \end{aligned}$$

in the update scheme

$$\mathbf{A}[\mathbf{x}]\mathbf{b}[\mathbf{xs}]\mathbf{c}[\mathbf{ys}]\mathbf{D}[\mathbf{zs}]\mathbf{e}\Longrightarrow.$$

\mathbf{A} and \mathbf{D} are ground, by virtue of being constants; \mathbf{b} is also ground, being equal to $\mathbf{A} + 1$; \mathbf{c} is semi-ground ($\mathbf{c} \in \{\mathbf{A} + n \mid 1 \leq n \leq \mathbf{D} - \mathbf{A}\}$); and \mathbf{e} is non-ground, since the length of the object represented by \mathbf{zs} cannot be determined. Similarly \mathbf{x} is ground, \mathbf{xs} and \mathbf{ys} are semi-ground, and \mathbf{zs} is non-ground.

Semi-ground terms are defined by the following rules (rule 5 is explained on page 43, rule 7 in chapter seven):

1. A ground term is semi-ground.
2. If α and β are semi-ground, and $\alpha[\gamma]\beta$ is a locator expression then:
 - if $\alpha[\gamma]\beta$ is on the left hand side of the update scheme, γ is semi-ground;
 - if $\alpha[\gamma]\beta$ occurs only in the right hand side of the update scheme, $|\gamma|$ is semi-ground ($|\cdot|$ is the length operator).

3. If α and $|\gamma|$ are semi-ground then so is β in the expressions $\alpha[\gamma]\beta$ and $\beta[\gamma]\alpha$.
4. The tuple $(\alpha_1, \dots, \alpha_n)$ is semi-ground if and only if $\alpha_1, \dots, \alpha_n$ are semi-ground.
5. If f is a well defined mapping for objects of type τ , and α is of type τ , then $f \alpha$ is semi-ground.
6. If the language of the nonterminal from the type grammar associated with an object γ is finite then $|\gamma|$ is semi-ground. If the language contains only one element $|\gamma|$ is ground.
7. If $m(\dots)$ is an archetype call, and the expressions appearing at the parameter positions in some subset p of the archetype's parameter position set are all semi-ground, then so are all expressions appearing at parameter positions in $\mathcal{G}_m(p)$, where \mathcal{G}_m is the grounding function of m . See chapter seven for definitions of 'parameter position set', and \mathcal{G}_m .

Note that no use is made of the guard in determining grounding. Any semi-ground term can be instantiated without reference to the guard.

Rules 1 to 3 derive grounding information from the structure of the update scheme under consideration. The asymmetry in rule 2 is due to the fact that a value can be found for a γ on the left hand side by examination of the current configuration, while this cannot be done for a γ which does not appear on the left hand side. All mappings involved in rule 5 are considered to have one argument, and rule 4 is provided for tupling and untupling arguments of mappings. In rule 5 a well defined mapping from objects of type τ to objects of type τ' is one in which every object of type τ' maps to a finite number of objects of type τ . The set of well defined mappings available may differ from implementation to implementation. An implementation specification must detail which mappings are defined. These must include the basic arithmetic operators, the length operator, and the concatenation operator and its inverse. It is by application of this last that `c` in example 2 is semi-ground, as shown in figure 1. Rule 6 makes use of typing information to determine whether an object has a finite number of possible lengths. Rule 7 is provided for completeness, the terminology used will be explained in chapter seven.

In contrast to Meijer's definition of traceability [53] the grounding rules above preclude the use of C-like strings, in which the end of the sequence is indicated by some special value. Such structures can however easily be defined by way of the archetype mechanism, introduced in chapter seven. Anticipating chapter seven, such a structure could be defined by

```
string([])    =  0|.
string(c s)  =  c string(s) | = [ c ≠ 0 ] ⇒ .
```

4 Memory Structures

The typing system also makes it possible to detect certain common types of addressing structure. This is achieved by means of a store-by-store

semi-ground terms	rule	newly semi-ground terms	mapping
—	(1)	A, D	
—	(6)	x	
A, x	(3)	b	
A, b	(2)	x	
b, D	(2)	xs ys	
xs ys	(5)	xs, ys	++ ⁻¹
xs, ys	(5)	xs , ys	·
b, xs	(3)	c	
b, c	(2)	xs	
c, D	(2)	ys	

Figure 1: Grounding terms in example 2

analysis of the update schemes in the update plan, and the way in which they are addressed. The aim of the analysis is to detect stacks, and variants and combinations of stacks, which define other common types of memory structure. Knowledge of addressing structures can be useful in determining semantic properties of update plans, as demonstrated in chapter six. The classification discussed below is summarised in figures 1 and 2 in appendix IV.

4.1 Structured Access

In the context of this analysis a locator expression will be called an *access* of the store in which it occurs. An access on the left hand side of an update scheme will be called a *read access*, and on the right hand side a *write access*. An access is *direct* if one of its locators is contained in a register. The register is said to address the access.

If a direct access occurs as part of one of the two patterns

$$R[l] \ l[x]r \implies R[l'] \ l'[y]r.$$

or

$$R[r] \ l[x]r \implies R[r'] \ l[y]r'.$$

then the access is said to be structured. In the first case the access is said to *read to the right* and *write to the left*; in the second the access *reads to the left* and *writes to the right*. In both cases, either of the accesses may be empty.

A register R is said to be *dedicated* to a store if its only use in an update plan is to address accesses of that store. A dedicated register may be *read dedicated* if it only addresses read accesses, and *write dedicated* if it only addresses write accesses. The locator in a dedicated register is said to be a *pointer* to the store it addresses.

A store is (read, write) structured if all (read, write) accesses of that store are consistently structured following one of the two patterns above, and each of these accesses has the same dedicated register. A store may be read *and* write structured without being (fully) structured. Only if

the same register is dedicated to read and write accesses is the store fully structured. A store which is read or write structured, but not fully structured, is said to be *semi-structured*.

If a store is read (write) structured, and reads (writes) to the left, those cells to the right of the read (write) pointer are called *lee* cells, cells to the left of the pointer are called *luff* cells. Similarly, if the store reads (writes) to the right cells to the left of the pointer are lee cells, cells to the right luff cells. Lee cells are those cells that have already been read (or written), luff cells have yet to be read (written). In a fully structured store lee and luff cells are determined by the direction in which the store is written.

4.2 Structured Stores

Stacks. A fully structured store is a stack.

Streams. Simpler versions of stacks are those in which there are either no write accesses or no read accesses — input and output streams respectively. These input and output streams are a slight generalisation of those introduced using the ‘?’ and ‘!’ notation in chapter two in that ‘?’ was only used for input streams which read to the right. Similarly, ‘!’ was only used for streams which wrote to the right.

4.3 Semi-Structured Stores

Queues. If a store is read and write structured, but not fully structured, it is a *queue*, if it satisfies the following conditions.

- It must read and write in the same direction.
- The read pointer may never point to a luff cell of the write pointer.

The second condition is probably only easy to determine if it is explicitly enforced by a guard at every read access.

Heaps. A simple form of heap is one in which contiguous sections of memory are allocated, and possibly written to, after which any cell in allocated memory may be read. This structure is reflected in the following conditions.

- The store must be write structured.
- Read accesses may only occur in the lee of the write pointer.

The latter condition is guaranteed if, with the exception of the write pointer, any object of the type of the heap, appearing on the right hand side of an update scheme, also appears on the left hand side.

A less restrictive definition of a heap is also possible. This requires that there be a register write dedicated to the heap, and that all other accesses conform to the second condition above.

4.4 Unstructured Stores

Read Only Memory. A read only, or static, store is one in which there are no write accesses.

Write Only Memory. The complement to a read only store is a write only store. In the context of an update plan this could represent, for example, some background storage device, or a graphical display device.

General Purpose Memory. A completely unstructured store represents the most general form of memory use. It is very difficult to derive semantic properties for this type of store.

5 Programme Stores

A special case of read-only memory is a *programme store*. This is the read only memory accessed by PC in a command driven plan. Note that the store addressed by PC in a command driven store *need not* be read only, and that a programme store therefore need not necessarily be read-only — there is no reason a machine with self-modifying code could not be specified. In the remainder of this thesis, however, all command driven update plans presented will have read-only programme stores, since this seems the most natural style of specifying instruction sets. Extending the relevant results to writeable programme stores should not present unsurmountable difficulties.

If, for a given command driven update plan, the register PC always addresses the same store, and this store is read only, then the store involved is said to be a *programme store*. The contents of a programme store are said to form a *programme*. A programme, together with the register PC and its contents, is said to form the *programme state*.

The semantic properties of programme stores will be analysed in chapter six.

6 Conclusions

This chapter primarily introduces concepts that will be of use in later chapters. These concepts are based on the typing mechanism. The two major concepts are grounding, which is relevant to the archetype mechanism introduced in chapter seven, and the classification of memory use paradigms, the semantic properties of which are discussed in chapter six, and applied in chapter eight.

An obvious question, when confronted with two update scripts, is to what extent the two scripts are “equivalent”, or whether one script is a “translation” of the other. In this chapter these concepts will be formalised, paving the way for their application in chapter eight.

Section 1 defines semantic equivalence, with varying degrees of freedom, between update plans. A concise diagrammatical notation for these definitions is introduced in section 2, and applied in section 4, in which some semantic properties of configurations with respect to some given update plan are defined. Section 5 shows how to use the properties defined in section 4, and the more restrictive equivalence properties from section 1, to prove more general equivalence. This method of proof is applied again in chapter eight.

1 Semantic Equivalence

In this section various semantic relations between update plans will be defined. The properties fall into two classes, *equivalences* and *translations*. These classes can be subdivided pairwise along three axes, as suggested in figure 1, with the default values being at the corner where the axes meet — e.g. ‘partial translation’ stands for ‘general transitive partial translation’, and ‘restricted local equivalence’ for ‘restricted full local equivalence’. The axes are characterised as follows (a fuller definition is

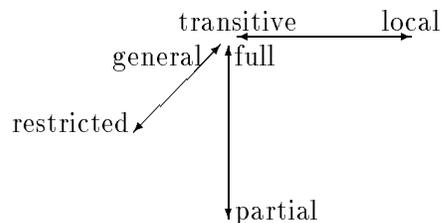


Figure 1: Types of semantic equivalence

given below):

full \leftrightarrow **partial** A semantic property is full if it applies to the whole of any configuration, and partial if it only applies to some well defined subset, for example some set of stores within configurations.

transitive \leftrightarrow **local** Local properties need only hold across one update, while transitive properties are, naturally, transitive.

general \leftrightarrow **restricted** A property is general if no restrictions (other than those implicit in its further classification) are imposed, and restricted if there *are* additional restrictions.

Equivalence and translation are defined in section 1.1. Section 1.2 introduces the full and partial variants of these properties, and section 1.3 the local and transitive forms. General and restricted properties will not be discussed until section 3, after the introduction of box diagrams in section 2.

1.1 Equivalences and Translations

In the context of these properties two update plans are *equivalent* if they have the same effect on configurations, i.e. if any development of one plan can be reached by the other, and vice versa. Update plans p_1 and p_2 are equivalent if, for any configuration c , it is true that $(c \Rightarrow_{p_1}^* c') \iff (c \Rightarrow_{p_2}^* c')$. The c 's in this relation need only be identical up to isomorphism. Any such isomorphism is assumed given, and is implicit in all definitions of (semantic) equivalence. The (semantic) equivalence symbol is ' \equiv '.

An update plan is a translation of another if the first can reach any configuration that the second can. The reverse need not be true. Plan p_1 is a translation of p_2 if

$$(c \Rightarrow_{p_1}^* c_1) \implies \exists c'[(c_1 \Rightarrow_{p_1}^* c') \wedge (c \Rightarrow_{p_2}^* c')]$$

and

$$(c \Rightarrow_{p_2}^* c') \implies (c \Rightarrow_{p_1}^* c'),$$

again for any configuration c . The translation symbol is ' \triangleright ', thus, above, $p_1 \triangleright p_2$. The terminology is inspired by compilers of high level languages in which one high level construct is translated into many low level constructs. The low level translation may pass through many configurations invisible at the higher level.

Where, in the definition of equivalence, two configurations were considered equal modulo isomorphism, in translation two configurations are "equal" modulo a homomorphism from configurations on the right hand side of the translation (in this case p_2) to configurations on the left hand side (p_1). This implies that the definition is only required to hold for configurations of p_1 that are images of configurations of p_2 . (A configuration is said to be a configuration of update plan p if it is intended for inclusion with p in an update script.)

1.2 Full and Partial Properties

Often a specification will include “important” and “unimportant” stores. Sometimes a (possibly dynamically determined) part of an individual store may be “important”, and the rest “unimportant”. It is often useful to restrict consideration of semantic equivalence to relevant sets of (portions of) stores. In partial equivalence (or translation), two developments c_1 and c_2 are considered to be the “same” if, for some set C of “important” cells, they can be written as $c \cup l_1$ and $c \cup l_2$ respectively, again modulo an iso- or homomorphism, with $\text{Dom } c \subseteq C$ and $\text{Dom } l_i \cap C = \emptyset$, for $i \in \{1, 2\}$. The update plans are said to be equivalent *within* (or to be a translation within) C . In discussions involving both partial properties and semantic irrelevance (defined in section 4) it will be useful to be able to refer to the sets of “unimportant” stores. The terminology *outside* will be used as the converse of ‘within’ — i.e. if two plans are equivalent within a set C , they can also be said to be equivalent outside the complement of C .

Partiality is indicated by superscripting the equivalence or translation symbol with the set within which the property holds, e.g. $p_1 \equiv^C p_2$.

1.3 Transitive and Local Properties

Any of the above properties is local if it holds for configurations derived in zero or one update, but not necessarily for configurations derived in any number of steps. For example, full local equivalence is defined by $(c \Rightarrow_{p_1}^{(0|1)} c') \iff (c \Rightarrow_{p_2}^{(0|1)} c')$, (where $c \Rightarrow_{p_1}^{(0|1)} c'$ is equivalent to $c \Rightarrow c' \vee c = c'$) and full local translation by

$$\begin{aligned} & (c \Rightarrow_{p_1}^{(0|1)} c_1) \implies \exists c' [(c_1 \Rightarrow_{p_1}^* c') \wedge (c \Rightarrow_{p_2}^{(0|1)} c')] \\ \wedge \\ & (c \Rightarrow_{p_2} c') \implies (c \Rightarrow_{p_1}^* c'). \end{aligned}$$

Locality is indicated by subscripting the equivalence or translation symbol with a ‘1’.

It is obviously far easier to prove local properties than transitive properties. Section 4 defines new properties that can be used to derive transitive properties from their local counterparts.

2 Box Diagrams

The properties above all take the form of existential dependencies. This section presents a diagrammatic notation, known as *box diagrams*, for such dependencies. Existence of any part of the “boxed” part of a box diagram guarantees the existence of the whole diagram. Here ‘part’ means a set of objects (configurations) from the box diagram, plus any relations specified between the elements of the set. The part of the diagram from which the existence of the remainder of the diagram is deduced is called the *seed*.

For example, the full transitive translation property, $p_1 \triangleright p_2$, can be expressed by the following box diagram.

$$\boxed{\begin{array}{c} c \xrightarrow{p_1}^* c_1 \\ p_2 \Downarrow^* \\ c' \end{array}} \xrightarrow{p_1}^* c'$$

The seeds in this box diagram are $\{c\}$, $\{c_1\}$, $\{c'\}$, $\{c, c_1 \mid c \Rightarrow_{p_1}^* c_1\}$, $\{c, c' \mid c \Rightarrow_{p_2}^* c'\}$, $\{c_1, c'\}$ and $\{c, c_1, c' \mid c \Rightarrow_{p_1}^* c_1 \wedge c \Rightarrow_{p_2}^* c'\}$. The essential existential dependencies are those already given in section 1.3, namely

$$\begin{aligned} & \forall c, c_1 [(c \Rightarrow_{p_1}^* c_1) \implies \exists c' [(c_1 \Rightarrow_{p_1}^* c') \wedge (c \Rightarrow_{p_2}^* c')]] \\ & \wedge \\ & \forall c, c' [(c \Rightarrow_{p_2}^* c') \implies \exists c_1 [c \Rightarrow_{p_1}^* c_1 \Rightarrow_{p_1}^* c']] \end{aligned}$$

but it also implies the trivial dependencies (the first, for example, is satisfied by $c_1 = c' = c$)

$$\begin{aligned} & \forall c [\exists c_1, c' [(c \Rightarrow_{p_1}^* c_1 \Rightarrow_{p_1}^* c') \wedge (c \Rightarrow_{p_2}^* c')]] \\ & \wedge \\ & \forall c_1 [\exists c, c' [(c \Rightarrow_{p_1}^* c_1 \Rightarrow_{p_1}^* c') \wedge (c \Rightarrow_{p_2}^* c')]] \\ & \wedge \\ & \forall c' [\exists c, c_1 [(c \Rightarrow_{p_1}^* c_1 \Rightarrow_{p_1}^* c') \wedge (c \Rightarrow_{p_2}^* c')]] \\ & \wedge \\ & \forall c_1, c' [\exists c [(c \Rightarrow_{p_1}^* c_1 \Rightarrow_{p_1}^* c') \wedge (c \Rightarrow_{p_2}^* c')]] \\ & \wedge \\ & \forall c, c_1, c' [((c \Rightarrow_{p_1}^* c_1) \wedge (c \Rightarrow_{p_2}^* c')) \implies (c_1 \Rightarrow_{p_1}^* c')] \end{aligned}$$

It is no accident that the essence of the box diagram is that part generated by non-trivial seeds — i.e. seeds which themselves contain some restriction on their constituent configurations. When reading box diagrams it is advisable to first consider such seeds.

Box diagrams can be used for a simple form of proof by diagram chasing. Any box diagrams with overlapping boxed parts may be combined to form new box diagrams. See section 4.1 for a simple example of diagram chasing, and section 5 for a more complex example.

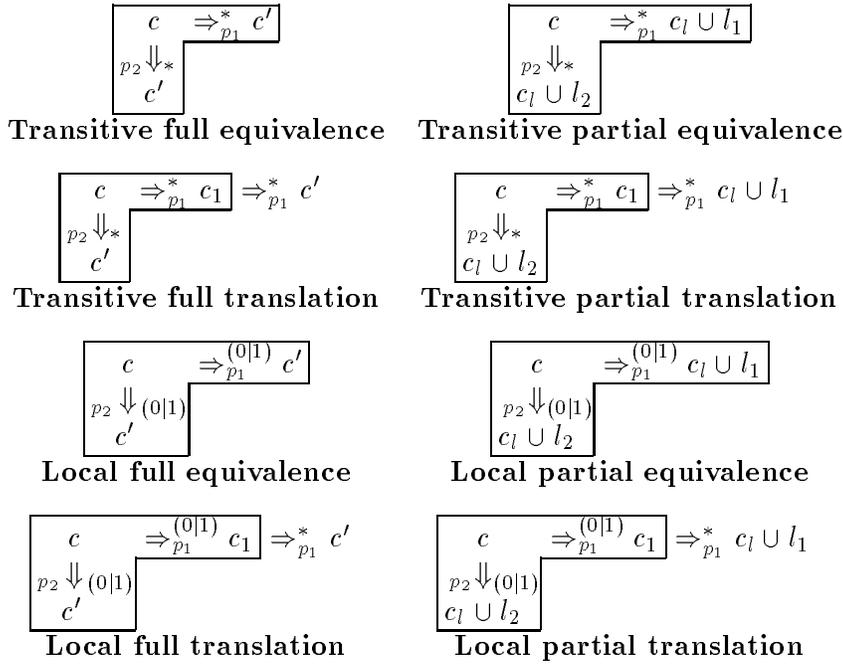
The eight properties given by the combination of the “equivalence \leftrightarrow translation” dichotomy and the “local \leftrightarrow transitive” and “full \leftrightarrow partial” axes in figure 1 are defined by the box diagrams in figure 2.

3 Restricted and General Properties

The third axis in figure 1 covers the possibility of adding extra conditions to a box diagram. Such extra restrictions apply to objects occurring within the box part of a diagram. If any such object is part of a seed the relevant restriction(s) must be included in the preconditions for the existence of the rest of the diagram. A typical restriction might require, for example, a stack pointer not to point to the bottom of a stack.

4 Semantic Irrelevance

The converse of the partiality above is the concept of *semantic irrelevance*. The partiality of partial equivalence (or translation) is a statement that part of the configuration is of no interest in a (final) development, though it may have had an effect on the derivation of that development. A typical example, in the specification of some machine, would be a set of temporary registers — the contents of temporary registers are almost certainly relevant to the semantics at intermediate stages, but of no interest once a final development is reached. Semantic irrelevance identifies



If the partiality of the above properties is outside L then, in all diagrams, $\text{Dom } c_l \cap L = \emptyset$ and $\text{Dom } l_i \subseteq L$ for $i \in \{1, 2\}$.

Figure 2: Semantic properties of update plans

part of the configuration as having no effect on the semantics with respect to the rest of the configuration, although it may well be of interest. An output stream is a good example of a semantically irrelevant store — an output stream does not affect the development of an update plan, but output is almost certainly a feature of interest in the semantics.

While equivalence and translation are relations between update plans, irrelevance is a property of sets of locators relative to a given plan.

4.1 Irrelevance

For some update plan p , the set of locators I is semantically irrelevant to configuration c if the contents of the cells addressed by locators in I are irrelevant to developments of c , i.e.

$$\boxed{\begin{array}{l} c_i \cup i_1 \Rightarrow^* c'_i \\ c_i \cup i_2 \Rightarrow^* c'_i \end{array}}$$

where $\text{Dom } i_j \subseteq I$ for $j \in \{1, 2\}$ and $\text{Dom } c_i \cap I = \text{Dom } c'_i \cap I = \emptyset$. The existence of either of the derivations in the above box diagram guarantees the existence of the other. Again local, partial and restricted versions can be constructed. Local partial irrelevance is, for example, defined by

$$\boxed{\begin{array}{l} c_i \cup i_1 \Rightarrow^{(0|1)} c_l \cup l_1 \\ c_i \cup i_2 \Rightarrow^{(0|1)} c_l \cup l_2 \end{array}}$$

with $\text{Dom } c_i \cap I = \emptyset$, $\text{Dom } c_l \cap L = \emptyset$, and $\text{Dom } l_j \subseteq L$ for $j \in \{1, 2\}$. L will typically be a subset of I .

4.2 Store Structure and Semantic Irrelevance

It is easy to prove the following facts. (See chapter five for a definition of lee and luff cells.)

- a write-only store is semantically irrelevant
- lee cells in a read-only store are semantically irrelevant
- luff cells (with respect to the write pointer) are semantically irrelevant in queues, stacks and heaps

The last case is a special case of the general rule that, if it is possible to prove for a write structured store that a read access will always take place in the lee of the write pointer, then the luff cells of that write pointer are semantically irrelevant.

Note that the first property is general (a write-only store is a write-only store, irrespective of the current configuration), while the remaining properties are restricted — precisely which cells are lee or luff cells depends on the current configuration.

4.3 Diagram Chasing

In section 5.1 transitive translation properties will be derived from local translation and local and transitive irrelevance by diagram chasing. As preparation for this proof, a simple example of proof by (box) diagram chasing will be presented. It will be shown how to establish transitive irrelevance properties from their local counterparts.

Take, for example, transitive partial irrelevance. This can be derived from local partial irrelevance by inductive diagram chasing. Let I be the irrelevant set, and L be the set outside of which the irrelevance holds. Assume furthermore that L is a subset of I . The local irrelevance property is then expressed by

$$\boxed{\begin{array}{l} c_i \cup i_1 \Rightarrow^{(0|1)} c_l \cup l_1 \\ c_i \cup i_2 \Rightarrow^{(0|1)} c_l \cup l_2 \end{array}}$$

with $\text{Dom } c_l \cap L = \emptyset$ and $\text{Dom } l_j \subseteq L$ for $j \in \{1, 2\}$. Since L is a subset of I , $c_l \cup l_1$ can be rewritten as $c'_i \cup i'_1$ with $c'_i \cap I = \emptyset$ and $i'_1 \subseteq I$, by “transferring” part of c_l to i'_1 . Similarly $c_l \cup l_2$ can be rewritten as $c'_i \cup i'_2$, and the local irrelevance property can be applied again, giving

$$\boxed{\begin{array}{l} c_i \cup i_1 \Rightarrow^{(0|1)} c'_i \cup i'_1 \Rightarrow^{(0|1)} c''_i \cup l'_1 \\ c_i \cup i_2 \Rightarrow^{(0|1)} c'_i \cup i'_2 \Rightarrow^{(0|1)} c''_i \cup l'_2 \end{array}}$$

The diagram can be rewritten as

$$\boxed{\begin{array}{l} c \Rightarrow^{(0|1|2)} c'_i \cup l'_1 \\ c \Rightarrow^{(0|1|2)} c'_i \cup l'_2 \end{array}}$$

and states that for any l'_2 , if $c'_i \cup l'_1$ can be derived from c in 0,1 or 2 steps $c'_i \cup l'_2$ can also be derived from c in 0,1 or 2 steps, and vice versa. Proofs

for other flavours of irrelevance are similar. Transitive irrelevance does not necessarily imply local irrelevance.

5 Deriving Transitive from Local Properties

Diagram chasing can be used not only to derive transitive irrelevance properties, but also, given an appropriate irrelevance property, to derive transitive translation (or equivalence) from the corresponding local property.

5.1 Methodology

The methodology will again be illustrated by an example. Local partial (outside some set of locators L) translation is defined by:

$$\begin{array}{ccc} c & \Rightarrow_{p_1}^{(0|1)} c_1 & \Rightarrow_{p_1}^* c_l \cup l_1 \\ p_2 \Downarrow (0|1) & & \\ c_l \cup l_2 & & \end{array}$$

Possibly $c_l \cup l_2$ can be developed further, $c_l \cup l_2 \Rightarrow_{p_2}^{(0|1)} c'_l \cup l'_2$, say. Local irrelevance gives:

$$\begin{array}{cc} c_l \cup l_2 & c_l \cup l \\ p_2 \Downarrow (0|1) & p_2 \Downarrow (0|1) \\ c'_l \cup l'_2 & c'_l \cup l'_2 \end{array}$$

(The diagram has been rotated 90° with respect to the corresponding diagram in section 4.) These two diagrams can be combined to give:

$$\begin{array}{ccc} c & \Rightarrow_{p_1}^{(0|1)} c_1 & \Rightarrow_{p_1}^* c_l \cup l_1 \\ p_2 \Downarrow (0|1) & & \\ c_l \cup l_2 & & c_l \cup l \\ p_2 \Downarrow (0|1) & & p_2 \Downarrow (0|1) \\ c'_l \cup l'_2 & & c'_l \cup l'_2 \end{array}$$

Local partial translation then establishes:

$$\begin{array}{ccc} c_l \cup l & \Rightarrow_{p_1}^{(0|1)} c'_1 & \Rightarrow_{p_1}^* c'_l \cup l'_1 \\ p_2 \Downarrow (0|1) & & \\ c'_l \cup l'_2 & & \end{array}$$

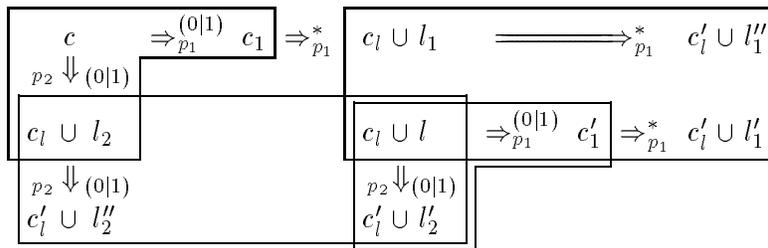
giving:

$$\begin{array}{ccc} c & \Rightarrow_{p_1}^{(0|1)} c_1 & \Rightarrow_{p_1}^* c_l \cup l_1 \\ p_2 \Downarrow (0|1) & & \\ c_l \cup l_2 & & c_l \cup l \Rightarrow_{p_1}^{(0|1)} c'_1 \Rightarrow_{p_1}^* c'_l \cup l'_1 \\ p_2 \Downarrow (0|1) & & p_2 \Downarrow (0|1) \\ c'_l \cup l'_2 & & c'_l \cup l'_2 \end{array}$$

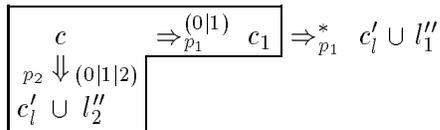
Finally transitive partial irrelevance can be written as:

$$\begin{array}{ccc} c_l \cup l_1 & \xRightarrow{p_1}^* c'_l \cup l'_1 \\ c_l \cup l & \Rightarrow_{p_1}^{(0|1)} c'_1 \Rightarrow_{p_1}^* c'_l \cup l'_1 \end{array}$$

and combined with the other diagrams to give:



A similar reasoning can be followed given a development $c_i \cup l_1 \Rightarrow_{p_1}^* c'_i \cup l''_1$ and it can therefore be concluded that:



thus establishing transitive partial translation from local partial translation and local and transitive partial irrelevance. Since local partial irrelevance implies transitive partial irrelevance, local partial translation and local partial irrelevance suffice to establish transitive partial translation.

The proofs for other flavours are again similar though, again, some boundary conditions may need to be checked.

5.2 Application

The semantic properties, and the proof method, described in this chapter can be applied to proving correctness of programme transformations.

For example, given two command driven update plans, and a “translation” from configurations for the one plan to configurations for the other (usually specifying transformations of the programme store and, implicitly, the contents of the PC, and possibly also of some ancillary stores and associated registers, such as a stack and its stack pointer) it is often the question whether the “translation” is a translation in the sense defined above.

It is usually relatively simple to prove local (partial) translation, possibly analysing possible configurations and proving restricted translation for separate classes of configuration. It is also usually not too difficult to show (or disprove) irrelevance of those parts of the configuration in which the translation is partial. The construction in section 5.1 can then be used to establish transitive translation. A similar method is applied in chapter eight.

6 Conclusions

This chapter has defined some semantic properties of update plans and scripts, and shown how to derive more general properties from more restricted ones. An initial description of an area of application of the subject matter of this chapter has been given. The methods of this chapter will be applied again in chapter eight.

Update Plans constitute a high level language for the specification of low level activities. They combine concrete and detailed low-level descriptions of data structures with an abstract and high-level algorithmic specification formalism. This chapter adds a potent macro-like mechanism to Update Plans, known as *archetypes*, which increases the expressive power of the formalism by facilitating abstraction away from unnecessary detail, and making it possible to specify concrete machines in an easy and natural way. The archetype mechanism is also eminently suited for the specification of parallel architectures.

Archetypes resemble parameterised macros and represent substructures of update schemes, making a form of structured programming possible. The parameter mechanism is similar to that of attribute grammars[18, 23, 39, 40, 44, 49], with parameters being either inherited or derived (synthesised).

1 Archetypes

Update Plans have already been used for the specification of abstract machines. The complexity of such machines is usually to be found in their command set, rather than in their addressing modes. Each command tends to have only a small number of addressing modes associated with it — in most cases the instruction determines the addressing mode uniquely. Instructions on concrete machines, on the other hand, are often formed by combining elements of two more or less orthogonal sets: the commands and the addressing modes. Specifying each possible combination of command and addressing mode leads to an explosive growth in the number of update schemes needed.

The archetype mechanism greatly increases the expressive power of Update Plans. Using the archetype mechanism complicated pointer structures, families of such structures, and even infinite classes of arbitrarily large structures may be replaced by a single archetype call, thus making it possible to express many update schemes as one.

Archetypes are inspired by macro mechanisms. Their parameter resolution system is purely “macro” in flavour, though their expansion, in particular that of recursive archetypes, may be context driven, i.e. dependent on the configuration in which the macro is expanded. Archetypes bear a degree of similarity to *syntax macros* [46].

Archetypes may be used both in update schemes and in specifications of configurations. The expansion mechanism is slightly more complicated for archetype applications in update schemes, due to the more complex structure of update schemes. In order to give a full presentation the expansion mechanism for archetype applications in update schemes will therefore be given. An archetype application in a configuration is expanded by ignoring any guards and right hand sides in its definition(s), and expanding it as if it were a *left-handed* archetype (see page 59).

1.1 Targeted Use

The primary application of the archetype mechanism is abstracting away the details of addressing modes.

Example 1

Given a suitable set of archetypes defining addressing modes, and the pointer structures associated with these modes, and specifying which of these may be used as source and destination operands (indicated here by `src` and `dst` respectively) the update scheme

$$\text{ADD } \text{src}(x) \text{ dst}(ea, y) \Rightarrow ea[x + y].$$

will expand to a number of schemes, one for each permissible combination of addressing modes. The source may, for example, be addressed in “register deferred” mode and the destination in “predecrement” mode in which case the relevant expansion of the above scheme could be

$$\begin{array}{l} \text{ADD REGDEF } i \text{ PREDEC } j. \\ \quad i[a] \ a[x] \ j[q] \ p[y]q \\ \Rightarrow \\ \quad \quad \quad j[p] \ p[x + y]. \end{array}$$

This example anticipates some of the syntactic sugar to be introduced in section 2.

A secondary application is the specification of operations on recursive data structures, as illustrated in example 6.

1.2 Syntax

Archetype Definitions. The syntax of archetype definitions is given by the following context free grammar. A guard may not contain an archetype call (this has to do with the structured nature of archetype calls), and a locator may not contain a call of a recursive macro (this has to do with the expansion mechanism).

$$\langle \text{item} \rangle \rightarrow \langle \text{archetype definition} \rangle$$

$$\langle \text{archetype definition} \rangle \rightarrow$$

$$\quad \langle \text{command archetype definition} \rangle \mid$$

$$\quad \langle \text{basic archetype definition} \rangle$$

$$\langle \text{basic archetype definition} \rangle \rightarrow$$

$$\quad \langle \text{basic declaration} \rangle \langle \text{basic definition} \rangle^+$$

$$\langle \text{basic declaration} \rangle \rightarrow \langle \text{basic archetype name} \rangle \langle \text{parameters} \rangle$$

$$\langle \text{basic definition} \rangle \rightarrow = \langle \text{basic body} \rangle .$$

$$\langle \text{basic body} \rangle \rightarrow$$

$$\quad \langle \text{configuration} \rangle \langle \text{guard} \rangle \langle \text{configuration} \rangle \mid$$

$$\quad \langle \text{repeat} \rangle \langle \text{guard} \rangle \langle \text{configuration} \rangle \mid$$

$$\quad \langle \text{configuration} \rangle \mid$$

$$\quad \langle \text{text} \rangle \mid \langle \text{context} \rangle \langle \text{guard} \rangle \langle \text{context} \rangle \mid$$

$$\quad \langle \text{text} \rangle \mid \langle \text{repeat} \rangle \langle \text{guard} \rangle \langle \text{context} \rangle \mid$$

$$\quad \langle \text{text} \rangle \mid \langle \text{context} \rangle$$

$$\langle \text{command archetype definition} \rangle \rightarrow$$

$$\quad \langle \text{command declaration} \rangle \langle \text{command definition} \rangle^+$$

$$\langle \text{command declaration} \rangle \rightarrow \langle \text{command archetype name} \rangle$$

$$\langle \text{parameters} \rangle \langle \text{text} \rangle$$

$$\langle \text{command body} \rangle \rightarrow = \langle \text{command body} \rangle .$$

$$\langle \text{command body} \rangle \rightarrow$$

$$\quad \langle \text{context} \rangle \langle \text{guard} \rangle \langle \text{configuration} \rangle \mid$$

$$\quad \langle \text{repeat} \rangle \langle \text{guard} \rangle \langle \text{configuration} \rangle \mid$$

$$\quad \langle \text{context} \rangle$$

$$\langle \text{parameters} \rangle \rightarrow (\{ \langle \text{term} \rangle , \}^*)$$

This grammar specifies the sugared syntax of archetype definitions. Syntactic sugar for archetypes is defined in section 2. A *command archetype name* is an upper case identifier, a *basic archetype name* a lower case identifier.

Archetype Calls. The syntax of an archetype call is basically that of a archetype declaration. There is one difference — archetype calls occur in pairs, and the calls may be “coupled” by an index. A term is now an expression built from constants, variables, archetype calls and operators. The following production rules should be added to the Update Plan grammar.

$\langle \text{term} \rangle \rightarrow \langle \text{archetype call} \rangle$

$\langle \text{archetype call} \rangle \rightarrow \langle \text{archetype name} \rangle \langle \text{index} \rangle\text{-opt} \langle \text{parameters} \rangle$

$\langle \text{archetype name} \rangle \rightarrow$
 $\langle \text{command archetype name} \rangle \mid$
 $\langle \text{basic archetype name} \rangle$

$\langle \text{index} \rangle \rightarrow \langle \text{number} \rangle \mid \{ \langle \text{number} \rangle \}$

An archetype's definition specifies text to be added to both the left and right hand sides of the update scheme in which it is called. As a consequence any archetype call on the left hand side of an update scheme must have a matching call on the right hand side, and vice versa. Since there may be more than one such pair of calls in some update scheme it is necessary to indicate which calls are coupled. This is done by indexing archetype call with a number, placed between the archetype's name and its parameter list. To distinguish between $\text{call}_1(\dots)$, an indexed call of the archetype $\text{call}(\dots)$, and $\text{call}_1(\dots)$, an unindexed call of $\text{call}_1(\dots)$, indices may be enclosed in braces ($\text{call}\{1\}(\dots)$). Each archetype-index pair must occur exactly once on both the left and right hand sides of the update scheme in which the archetype is called. The index may be omitted in some cases, as may one member of a pair of a calls, as detailed in section 2. In the following indices will be omitted, where this does not lead to confusion. Example 2 illustrates archetype expansion.

Expansion takes place, conceptually, as follows. For each pair of archetype calls, the archetype calls are replaced by the corresponding expansions, and the left and right hand side contexts are added to the corresponding sides of the update scheme in which the calls occurred. If more than one expansion is possible copies of the update scheme are generated for each possible expansion. See sections 3.2, 4.1 and 4.2 for more details.

————— Example 2 —————

Consider the archetype definition

$$\text{pop}(\mathbf{v}) = \mathbf{s} \ \mathbf{s}[\mathbf{v}]\mathbf{t} \Longrightarrow \mathbf{t}.$$

Textual expansion (see section 3.1) of this archetype in the update scheme

$$\text{ADD SP}[\text{pop}_1(\mathbf{x})] \ \text{ACC}[\mathbf{y}] \Longrightarrow \text{SP}[\text{pop}_1(\mathbf{x})] \ \text{ACC}[\mathbf{x} + \mathbf{y}].$$

will give

$$\text{ADD SP}[\mathbf{s}] \ \text{ACC}[\mathbf{y}] \ \mathbf{s}[\mathbf{v}]\mathbf{t} \Longrightarrow \text{SP}[\mathbf{t}] \ \text{ACC}[\mathbf{x} + \mathbf{y}].$$

The call of $\text{pop}(\mathbf{x})$ on the left hand side has been replaced by \mathbf{s} , and that on the right hand side by \mathbf{t} . The locator expression

$s[v]t$ has been added to the left hand side. Parameter resolution (see section 3.2) equates x to v giving

$$\text{ADD SP}[s] \text{ ACC}[y] \ s[v]t \Longrightarrow \text{SP}[t] \ \text{ACC}[v + y].$$

Again, this example anticipates some of the syntactic sugar in section 2.

2 Syntactic Sugar

The contexts of archetype definitions inherit syntactic sugar from standard update schemes, enabling unnecessary locators and guards to be omitted. In addition the following archetype specific syntactic sugar is defined. The terminology used is that of the grammar in section 1.2.

2.1 Redundancy in Calls

Left- and Right Handed Archetypes. If, for a given archetype, in all of its definitions, the expansion of the right hand side is empty then the right hand side call of that archetype may be omitted. If all right hand side calls are omitted such an archetype is called a *left handed* archetype. The same holds for empty left hand sides, in which case the archetype is *right handed*. In theory an archetype could be both left and right handed, in which case both calls would be omitted. Since the conditions for omission only guarantee that the expansions of the archetype are empty this could result in “invisible” archetype calls introducing supplementary locator expressions, and is therefore forbidden.

Context Independent Archetypes. Left and right handed archetypes have “singleton” — i.e. unpaired — calls. They are restricted in that a left handed archetype may only be called on the left hand side, and a right handed archetype on the right hand side. An *ambidextrous archetype* is a “singleton” archetype which does not have this restriction. In the definition of an ambidextrous archetype the expansion is separated from the contexts by a vertical line (see example 3). A definition of an ambidextrous archetype m is equivalent to a pair of archetype declarations, of m_l and m_r say, where m_l is a left handed archetype and m_r is right handed.

Example 3

The archetype definition

$$m(\dots) = \text{text} \mid \text{lhs} \text{ =}[\text{guard}] \Rightarrow \text{rhs}.$$

is equivalent to the two archetype definitions

$$m_l(\dots) = \text{text} \ \text{lhs} \text{ =}[\text{guard}] \Rightarrow \text{rhs}.$$

and

$$m_r(\dots) = \text{lhs} \text{ =}[\text{guard}] \Rightarrow \text{text} \ \text{rhs}.$$

where calls of m on the left hand side are equivalent to calls of m_l , and right hand side calls of m are equivalent to calls of m_r .

Since it cannot be determined from an archetype definition where an archetype call appearing in one of its parameters may be applied, only ambidextrous archetypes may appear in parameter expressions.

Omitting Indices. The sugar defined above produces archetypes with “singleton” calls — in contrast to standard archetypes, calls of which are always paired. Since there is no coupling, indexing is redundant and may be omitted. Similarly, if there is only one paired call of some archetype in an update scheme that archetype need not be indexed.

2.2 Redundancy in Definitions

Command Archetypes. If m is an ambidextrous archetype and the expansions of all definitions of m begin with the same constant then that constant may be used as the archetype name, if it has not already been so used. The remainder of the expansion is then placed between the archetype declaration and the archetype definition symbol. Such an archetype is called a *command archetype*. The command archetype

$$\text{CONST}(\dots) \text{ text} = \text{lhs} \text{ [= [guard]} \Rightarrow \text{rhs}.$$

is the sugared version of

$$m(\dots) = \text{CONST text} \mid \text{lhs} \text{ [= [guard]} \Rightarrow \text{rhs}.$$

with calls of m replaced throughout the update plan by calls of `CONST`.

Empty Guards and Right Hand Sides. If the guard and right hand side of an archetype definition are both empty then the transition symbol (\Rightarrow) may be omitted.

Shared Formals. Identical archetype declarations may be shared. The period is omitted from the end of each shared definition but the last. The ‘=’ is not omitted.

Don’t Care Values. Irrelevant parameters may be replaced by ‘_’, which is the “don’t care” symbol.

3 Expansion

The archetype expansion mechanism, inspired by traditional macro expansion, consists of two stages, the *textual expansion stage* and the *parameter resolution stage*, as described in sections 3.1 and 3.2 respectively. In the following the update scheme in which an archetype call occurs will be referred to as *the* scheme of that call. Archetype calls in an expansion or context of an archetype, or as parameter of an archetype, inherit their scheme from the caller.

3.1 Textual Expansion

Both the left and right hand sides of an archetype body consist of two parts, the *expansion* and the *context*. The expansion is the text that actually replaces the archetype call, the left hand side expansion replacing the call on the left hand side of the scheme, and the right hand side

expansion that on the right hand side. The contexts must be added to the call's scheme; again the left hand (right hand) side context is simply textually added to the left hand (right hand) side of the scheme. The guard of the call is joined to the guard of the scheme using conjunction.

More formally, if $(\lambda x.(lhs \equiv [guard] \Rightarrow rhs))(m(p_1, \dots, p_n))$ is an update scheme containing a call of the archetype $m(p_1, \dots, p_n)$, and an instantiation of this archetype is

$$m(p_1, \dots, p_n) = txt_l \ sch_l \equiv [grd] \Rightarrow txt_r \ sch_r.$$

then the result of textually expanding the archetype in the scheme is

$$\begin{aligned} lhs' & \equiv [guard \wedge grd] \Rightarrow rhs'. \\ \text{where } lhs' & = (\lambda x.lhs)txt_l \ sch_l \\ rhs' & = (\lambda x.rhs)txt_r \ sch_r \end{aligned}$$

The definition of textual expansion may seem unduly complicated, but the most obvious alternative, integral *in situ* expansion, while attractive in its simplicity, can lead to unexpected results, as shown in example 4.

Example 4

Consider the following archetype definition, using an ad hoc notation, designed for *in situ* expansion (note that the syntax of this “archetype” does not conform completely to the syntax for archetypes with separate expansion of the text and contexts)

$$\mathbf{ref}(a) = a[b] \ b \Longrightarrow .$$

When called as the left locator of a locator expression, using integral *in situ* expansion, the effect will be as suggested by the archetype's name. The archetype call

$$\dots \mathbf{ref}(a)[v] \dots \Longrightarrow \dots .$$

will expand to

$$\dots a[b] \ b[v] \dots \Longrightarrow \dots .$$

When called as the right locator, however, expansion gives unexpected results. For example

$$\dots [v]\mathbf{ref}(a) \dots \Longrightarrow \dots .$$

gives

$$\dots [v]a[b] \ b \dots \Longrightarrow \dots .$$

Even more problematic is the case of archetype calls within locator expressions.

3.2 Parameter Resolution

If the parameters of archetype calls and definitions were always simply variables parameters would not need to be resolved — the variables could simply be replaced throughout the scheme by the corresponding parameters from the archetype's definition. That a more complicated approach is needed if parameters may be compound expressions can be seen in the following example.

Example 5 (a)

Consider the archetype definition

$$\mathbf{bdisp}(b + d) = \mathbf{BDISP} \ r \ d \ r[b] \implies.$$

and its application in

$$\mathbf{CEQ} \ \mathbf{bdisp}_1(x) \ \mathbf{bdisp}_2(x) \implies \mathbf{CC}[\mathbf{TRUE}].$$

(Intuitively \mathbf{CEQ} is comparing the effective addresses of the two \mathbf{BDISP} operands. A more realistic example illustrating the problem could be given, but it would probably lack the simplicity of this slightly forced example.) Let

$$\mathbf{bdisp}(b_1 + d_1) = \mathbf{BDISP} \ r_1 \ d_1 \ r_1[b_1] \implies.$$

be the instantiation of the definition which is to be substituted for $\mathbf{bdisp}_1(x)$. Consistent substitution of $b_1 + d_1$ for x will give

$$\mathbf{CEQ} \ \mathbf{BDISP} \ r_1 \ d_1 \ \mathbf{bdisp}_2(b_1 + d_1) \ r_1[b_1] \implies \mathbf{CC}[\mathbf{TRUE}].$$

If $\mathbf{bdisp}_2(b_1 + d_1)$ is to be expanded according to the instantiation

$$\mathbf{bdisp}(b_2 + d_2) = \mathbf{BDISP} \ r_2 \ d_2 \ r_2[b_2] \implies.$$

an expression $(b_2 + d_2)$ is available which can be substituted for $b_1 + d_1$, but no expressions are available for b_1 and d_1 individually.

A parameter resolution method is needed for situations such as these. Rather than using a consistent substitution mechanism, parameters are resolved by maintaining a system of equations as rewriting takes place. It should be emphasised that parameter resolution remains a purely textual manipulation. In example 5(a) resolution will yield expressions for b_1 and b_2 , possibly rewriting other expressions in the update scheme, or adding equalities between expressions to the guard.

In this example, the first rewrite will give

$$\begin{array}{l} \text{CEQ BDISP } \mathbf{r}_1 \ \mathbf{d}_1 \ \mathbf{bdisp}_2(\mathbf{x}) \ \mathbf{r}_1[\mathbf{b}_1] \\ \implies \text{CC}[\text{TRUE}]. \end{array}$$

and the equation $x = b_1 + d_1$. The second archetype can now be rewritten giving

$$\begin{array}{l} \text{CEQ BDISP } \mathbf{r}_1 \ \mathbf{d}_1 \ \text{BDISP } \mathbf{r}_2 \ \mathbf{d}_2 \ \mathbf{r}_1[\mathbf{b}_1] \ \mathbf{r}_2[\mathbf{b}_2] \\ \implies \text{CC}[\text{TRUE}]. \end{array}$$

with equations $\{x = b_1 + d_1, x = b_2 + d_2\}$.

The basis for parameter resolution is grounding, as defined in chapter five. The textual expansion mechanism described in section 3.1 may leave some terms non-ground (e.g. \mathbf{x} in example 2). The aim of parameter resolution is to derive semi-ground expressions for non-ground terms. The base for resolution is the set of equations generated during textual expansion. This set of equations is called the *resolution set*.

Parameters are resolved iteratively. Conceptually, examination of the scheme of the call, after expansion, determines which expressions are semi-ground. Semi-ground expressions are then found for as yet unresolved variables by applying the equations in the resolution set. The expressions found are then substituted for occurrences of the corresponding variables both in the scheme and in the resolution set, and the process is repeated until all terms are semi-ground. If at any point in this process a non-trivial equation is derived relating two semi-ground expressions this is added to the scheme's guard. This process is slightly complicated by the existence of recursive archetypes, expansion of which is driven by the current configuration, requiring expansion and resolution to be interleaved. Resolution, however, is independent of the expansion mechanism. The parameter resolution mechanism presented here is certainly no more expensive than unification.

Continuing the example

$$\text{CEQ } \mathbf{bdisp}_1(\mathbf{x}) \ \mathbf{bdisp}_2(\mathbf{x}) \implies \text{CC}[\text{TRUE}].$$

may expand to

$$\begin{array}{l} \text{CEQ BDISP } \mathbf{r}_1 \ \mathbf{d}_1 \ \text{BDISP } \mathbf{r}_2 \ \mathbf{d}_2 \\ \quad \mathbf{r}_1[\mathbf{b}_1] \quad \mathbf{r}_2[\mathbf{b}_2] \\ \implies \text{CC}[\text{TRUE}]. \end{array}$$

with the associated resolution set (semi-ground terms are in bold font) $\{x = \mathbf{b}_1 + \mathbf{d}_1, x = \mathbf{b}_2 + \mathbf{d}_2\}$. Substituting $\mathbf{b}_2 + \mathbf{d}_2$

for x throughout the scheme and the resolution set gives $\{b_2 + d_2 = b_1 + d_1, b_2 + d_2 = b_2 + d_2\}$. The non-trivial equation is added to the scheme's guard, giving

$$\begin{array}{l} \text{CEQ BDISP } r_1 \ d_1 \ \text{BDISP } r_2 \ d_2 \\ \quad \quad \quad r_1[b_1] \quad \quad \quad r_2[b_2] \\ \Rightarrow [b_2 + d_2 = b_1 + d_1] \Rightarrow \text{CC}[\text{TRUE}]. \end{array}$$

4 Recursion

Update schemes are instantiated to update rules by matching them against the current configuration, possibly using type information to determine the number of cells occupied by the value of some variable. This allows variables to represent simple structures; atomic types, arrays, records, etc. It would be useful if there were a mechanism for representing more complicated structures, such as trees. Recursive archetypes provide such a mechanism, at the cost of increasing the complexity of the instantiation mechanism. Archetype definitions can be seen as (embellished) rules in a context free grammar (as defined in section 4.1). Recursive archetypes are therefore expanded during instantiation in a manner akin to grammar driven recognition of strings in context free languages.

Example 6

The recursive archetype definitions

$$\begin{array}{l} \text{tree}() = \text{LEAF } x \Rightarrow \text{LEAF } x. \\ \text{tree}() = \text{NODE } \text{tree}_1() \ \text{tree}_2() \\ \quad \quad \quad \Rightarrow \text{NODE } \text{tree}_1() \ \text{tree}_2(). \end{array}$$

define a binary tree structure. A possible application is in the definition of balancing operations as used for AVL trees. An update scheme for one of these, the “up right” operation, is

$$\begin{array}{l} \text{UPR NODE } \text{tree}_1() \ (\text{NODE } \text{tree}_{21}() \ \text{tree}_{22}()) \\ \quad \quad \quad \Rightarrow \text{NODE } (\text{NODE } \text{tree}_1() \ \text{tree}_{21}()) \ \text{tree}_{22}(). \end{array}$$

The parentheses around the archetype expansions have been added to improve legibility and are not an essential part of the update scheme.

Recursive archetypes represent infinite hierarchies of expansions and can therefore not be expanded using information based solely on the update plan in which they are defined. Expansion of recursive archetypes is driven by the current configuration. Some guarantee is needed, based on information derived solely from the update plan, that parameter resolution is finite (i.e. that the resolution set is finite) and complete (i.e. that a solution to the resolution set can be found, or shown not to exist).

4.1 Finite Resolution

Parameter resolution will be finite if textual expansion occurs in a finite number of steps. A sufficient condition for this can be given in terms of the *archetype grammar*, a context free grammar based on archetype definitions.

Since it is the left hand side of an update scheme that is matched with the current configuration only information from the left hand sides of archetype definitions may be applied to guarantee termination. Only the expansion of the left hand side is considered, in order to simplify the termination conditions, though extension to cover the whole left hand side may be feasible. Archetypes appearing in an archetype body elsewhere than the expansion of the left hand side may only be recursive if coupled to a call that *is* in the expansion of the left hand side. Archetypes appearing in parameters may not be recursive.

The basis for the archetype grammar is the set of typing rules. The typing rules assign to each variable in an update plan an element, either terminal or nonterminal, of the type grammar, introduced in chapter five. The type grammar defines all types appearing in an update plan. Of particular importance is the distinction between atomic types and sequence types. Objects of an atomic type are assumed to occupy a fixed number of cells, while those of a sequence type occupy a variable number of cells, possibly zero.

In the type grammar atomic types are terminals, and for each non-atomic type there is a nonterminal, the *type generator* of that type, which is the left hand side of a production rule generating exactly the set of possible terminal type expressions for that type. If this is a singleton set its element may be used as type generator.

The archetype grammar is formed by adding production rules, based on the archetype definitions, to the type grammar. Archetype names are added to the set of nonterminals, and any constants appearing in an archetype body to the set of terminals. New production rules are added to the grammar, derived from the archetype definitions by removing the parameter lists from the archetype declarations and deleting all but the left hand side expansion of the archetype bodies. Any expression not soluble at compile time, i.e. any expression containing one or more variables or archetype calls, is replaced by its type generator.

Example 7

Given the type alphabet $\{\mathbf{Loc}, \mathbf{Num}\}$, the type grammar

$$\langle \mathbf{LocNum} \rangle \rightarrow \mathbf{Loc} \mid \mathbf{Num}$$

and the typing rules $\mathbf{v} \rightarrow \langle \mathbf{LocNum} \rangle$, and $(\mathbf{adr} - \mathbf{reg}_1(\mathbf{r})) \rightarrow \mathbf{Loc}$, the archetype grammar corresponding to the archetype

definitions:

$$\begin{aligned} \text{op}(\text{reg}(\mathbf{r})) &= \text{REG } \text{val}_1(\mathbf{r}) \Rightarrow \text{val}_1(\mathbf{r}). \\ \text{op}(\mathbf{v}) &= \text{BDISP } \text{val}_1(\mathbf{r}) (\text{adr} - \text{reg}_1(\mathbf{r})) \text{ adr}[\mathbf{v}] \\ &\quad \Rightarrow \text{val}_1(\mathbf{r}) \text{ reg}_1(\mathbf{r}). \end{aligned}$$

$$\text{reg}(\mathbf{r}) = \mathbf{v} \mid (\mathbf{R} + \mathbf{r})[\mathbf{v}] \Rightarrow .$$

$$\text{val}(\mathbf{v}) = \mathbf{v} \Rightarrow .$$

$$\text{val}(\mathbf{v}) = \text{op}_1(\mathbf{v}) \Rightarrow \text{op}_1(\mathbf{v}).$$

is:

$$\begin{aligned} \langle \text{op} \rangle &\rightarrow \mathbf{REG} \langle \text{val} \rangle \mid \\ &\quad \mathbf{BDISP} \langle \text{val} \rangle \mathbf{Loc} \\ \langle \text{reg} \rangle &\rightarrow \langle \text{LocNum} \rangle \\ \langle \text{val} \rangle &\rightarrow \langle \text{LocNum} \rangle \mid \\ &\quad \langle \text{op} \rangle \\ \langle \text{LocNum} \rangle &\rightarrow \mathbf{Loc} \mid \mathbf{Num} \end{aligned}$$

Finite resolution is guaranteed if no derivations of the type $\alpha \Rightarrow^+ \alpha$ are possible in the archetype grammar. This ensures that as expansion proceeds the length of the replacement text increases. Since only a finite sequence of cells will be defined at the location of the archetype call only a finite number of expansions is possible. This is the reason recursive archetypes calls may only occur in a term specifying the contents of (a sequence of) cells. In a realistic implementation additional restraints may be imposed on the archetype grammar in order to ensure efficient expansion.

4.2 Complete Parameter Resolution

In chapter five grounding by way of archetype calls was defined in terms of the as yet undefined parameter position set and archetype grounding function. These notions will now be defined.

For any given archetype $m(p_1, \dots, p_n)$ the parameter position set P_m is defined to be $\{1, \dots, n\}$. The parameter position set indexes the archetype's parameters. Any archetype call specifying the contents of (a sequence of) cells — i.e. a call not occurring in a locator — also has the implicit parameters $[l, l]$, $[r$ and $r]$. These implicit parameters are the left and right locators of the left and right hand side call of the archetype, and therefore of the corresponding expansion in the archetype definition. The left locator on the left hand side is indicated by $[l$, the right locator on the left hand side by $l]$; $[r$ and $r]$ are interpreted analogously.

For any archetype definition m the grounding function $g : 2^{P_m} \mapsto 2^{P_m}$ can be defined such that $g(p_{in}) = p_{out}$ is equivalent to “if all parameters at positions in p_{in} are semi-ground, then so are all parameters at positions in p_{out} ”. The value of p_{out} can be determined by applying the grounding rules in chapter five.

If archetype `neg` has definition

$$\mathbf{neg}(\mathbf{v}) = \mathbf{v} \implies -\mathbf{v}.$$

then its parameter position set is $\{[l, l], [r, r], 1\}$. For this definition some values of g are, assuming \mathbf{v} has atomic type

$$\begin{aligned} g(\{r\}) &= \{[r, r]\} \\ g(\{l\}) &= \{[l, l], 1\} \\ g(\{[l, r]\}) &= P_{neg} \end{aligned}$$

The function g is extended to cover all definitions of the archetype, giving the archetype grounding function \mathcal{G}_m , by taking the intersection of g over all definitions.

$$\mathcal{G}_m(p) = \bigcap_{g \in G} g(p)$$

where $G = \{g \mid g \text{ is the grounding function of some definition of } m\}$.

For any set of archetype definitions the corresponding set of archetype grounding functions can be determined by standard closure techniques.

5 Some Examples

Sizeable applications of the archetype mechanism can be found in chapters eight and nine. In this section the programme counter convention introduced in chapter two will be shown to be closely related to the archetype mechanism, and will be defined in terms of it. This opens the door to extensions of the convention, and two of these are introduced.

5.1 Hiding the Programme Counter

It should now be clear that an update scheme in which the programme counter has been “hidden” on both the left and right hand sides is in fact the body of an archetype definition. The convention is then equivalent to prefacing each such scheme with, for example, ‘ $\mathbf{pc}() =$ ’, where \mathbf{pc} is some archetype name not used elsewhere in the update plan, and adding the update scheme

$$\mathbf{PC}[\mathbf{pc}] \ \mathbf{pc}[\mathbf{pc}()] \mathbf{qc} \implies \mathbf{PC}[\mathbf{pc}'] \ \mathbf{pc}'[\mathbf{pc}()] \mathbf{qc}.$$

to the update plan.

Indeed, the syntactic sugar in section 2 ensures that, if every update scheme in an update plan is written as a command, it suffices to prefix the plan with

$$\begin{aligned} &\mathbf{PC}[\mathbf{pc}] \ \mathbf{pc}[\mathbf{pc}()] \mathbf{qc} \implies \mathbf{PC}[\mathbf{pc}'] \ \mathbf{pc}'[\mathbf{pc}()] \mathbf{qc}. \\ &\mathbf{pc}() = \end{aligned}$$

5.2 Asynchronous Parallel Processors

It is also possible to define more than one programme counter, thus specifying asynchronous parallel processors with shared memory and identical instruction sets.

$$\begin{aligned} PC_1[pc] \quad pc[pc()]qc &\Longrightarrow PC_1[pc'] \quad pc'[pc()]qc. \\ &\vdots \\ PC_n[pc] \quad pc[pc()]qc &\Longrightarrow PC_n[pc'] \quad pc'[pc()]qc. \end{aligned}$$

5.3 Specifying the Instruction Cycle

Another possibility is “hiding” some other structure. A lower level specification, for example of the instruction cycle, might be more legible with the clock hidden.

$$CLOCK[tick()] \Longrightarrow CLOCK[tick()].$$

For a simplified PDP-11 like machine the instruction cycle can be considered to consist of four steps; fetching the instruction (IFETCH), fetching the source operand (SRC), fetching the destination operand and its address (DST), and performing the operation and writing the result to the destination (EXEC). These steps can then be specified by:

$$\begin{aligned} tick() &= \text{IFETCH } PC[pc] \quad pc[\text{command}]qc \Longrightarrow \text{SRC } IR[\text{command}] \quad PC[qc]. \\ &\quad \vdots \\ &= \text{SRC } IR + S[\text{IMM}] \quad PC[pc] \quad pc[x]qc \Longrightarrow \text{DST } RX[x] \quad PC[qc]. \\ &\quad \vdots \\ &= \text{DST } IR + D[\text{REG } r] \quad R + r[y] \Longrightarrow \text{EXEC } MAR[R + r] \quad RY[y]. \\ &\quad \vdots \\ &= \text{EXEC } IR[\text{ADD}] \quad RX[x] \quad RY[y] \quad MAR[\text{dst}] \Longrightarrow \text{IFETCH } \text{dst}[x + y]. \\ &\quad \vdots \end{aligned}$$

where IR is the instruction register, MAR is the memory address register, RX and RY are the right and left port of the ALU respectively, and S and D are the offsets of the source and destination addresses in the command word. The specification assumes that the registers occupy a contiguous section of memory, with R as the left locator of register 0. A more elegant formulation of the same specification would be:

Addressing Modes.

$$\begin{aligned} op(R + r, v) &= \text{REG } r \quad R + r[v]. \\ op(pc, v) &= \text{IMM } PC[pc] \quad pc[v]qc \Longrightarrow PC[qc]. \\ &\quad \vdots \end{aligned}$$

Operands.

$$\begin{aligned} \text{src}(val) &= op(-, val). \\ \text{dst}(ea, val) &= op(ea, val). \end{aligned}$$

Instructions.

$$\begin{aligned} \text{IFETCH}() &= \text{PC}[\text{pc}] \text{ pc}[\text{command}] \text{qc} && \implies \text{IR}[\text{command}] \text{PC}[\text{qc}]. \\ \text{SRC}() &= \text{IR} + \text{S}[\text{src}(\text{x})] && \implies \text{RX}[\text{x}]. \\ \text{DST}() &= \text{IR} + \text{D}[\text{dst}(\text{ea}, \text{y})] && \implies \text{RY}[\text{y}] \text{MAR}[\text{ea}]. \\ \\ \text{EXEC}() &= \text{IR}[\text{ADD}] \text{RX}[\text{x}] \text{RY}[\text{y}] \text{MAR}[\text{dst}] \implies \text{dst}[\text{x} + \text{y}]. \\ &\vdots \end{aligned}$$

Instruction Cycle.

$$\begin{aligned} \text{tick}() &= \text{IFETCH}() \implies \text{SRC}. \\ &= \text{SRC}() \implies \text{DST}. \\ &= \text{DST}() \implies \text{EXEC}. \\ &= \text{EXEC}() \implies \text{IFETCH}. \end{aligned}$$

The Clock.

$$\text{CLOCK}[\text{tick}()] \implies \text{CLOCK}[\text{tick}()].$$

6 Conclusions

Update Plans have already been shown useful as a specification formalism for abstract machines. The addition of a macro-like mechanism not only makes them suited to the specification of concrete machines, it also opens the door to applications in the specification of asynchronous parallel processes. Another extension of Update Plans to cover synchronous parallelism is presented in chapter nine.

Although the mechanism presented here is specific to Update Plans, the basic principle of having macro definitions and calls reflect the structure of the formalism in which they are defined could almost certainly be fruitfully applied to other languages and formalisms with well defined structures — for example, as has been suggested [55], the method used to determine grounding sets for archetypes could probably be applied to the problem of determining critical positions in predicates of Extended Affix Grammars [24, 42, 53, 74].

Code generation often generates intermediate code, before register allocation, with a tree-like structure, in which subtrees represent intermediate results in a calculation. This code must then be linearised — to instructions on some concrete machine — by passing intermediate results via registers, or via a stack.

With Update Plans both the concrete machine and an abstract machine for the intermediate code can be specified, using the same formalism. A transformation from intermediate code to concrete machine code can then be given. Using the properties introduced in chapter five this transformation can be shown to preserve semantics — i.e. to provide a correct translation.

In this chapter two related machines are specified:

- a linear code machine — a PDP-11 style machine, not actually extant, but similar enough to concrete machines to be used as an example,
- a machine for a tree-structured language, the language being typical of intermediate code emitted by a compiler.

The machines specified are based on machines described by Giegerich [28]. Transformations are defined from tree code to linear code (register allocation) and it is shown that these transformations define a translation, in the sense of chapter 6.

The aim of this chapter is two-fold. Firstly, the machine specifications serve as extensive examples of applications of the archetype mechanism introduced in chapter seven. In particular the tree machine illustrated in figures 1 to 8 in appendix VI provides a sizeable example of archetype definitions and their expansion. Secondly, the chapter suggests a way in which formal methods may be applied to a problem — showing register allocation to be correct — which has not previously been readily amenable to such an approach.

Section 1 presents the linear machine, which is summarised in figures 1 and 2 of appendix VI, and section 2 the tree machine, summarised in figures 1 and 3 of appendix VI. The scheme and archetype numbering in these, and other, sections is provided for ease of reference, as in chapter four, and does not form part of the schemes and archetypes. The transformations from tree code to linear code are given in section 3 and the semantic equivalence to the original code of the resulting code is shown in section 4. A short summary, and some suggestions for more extensive applications are given in section 5.

The transformation is not wholly realistic, since the concrete machine is assumed always to have a sufficient number of registers. This restriction is addressed in section 5.

The methods presented here can be of importance in developing provably correct code generators.

1 The Linear Code Machine

The linear code machine has eight addressing modes, two data transfer commands, four arithmetic and six comparison operators, and five programme flow commands. At this level of specification the number of registers is considered to be infinite.

1.1 The Specification

Addressing Modes. The following archetypes all define two results, the effective address and the value. In all of these the locator contained in register \mathbf{r} , here consistently represented by \mathbf{b} , is called the *base value* of the addressing mode. The addressing modes and corresponding archetypes are:

register, $\text{REG}(\cdot, \cdot) \mathbf{r}$. Direct addressing, the value accessed is the value in the argument (register \mathbf{r}) of the addressing mode.

register deferred, $\text{REGDEF}(\cdot, \cdot) \mathbf{r}$. The address of the value is in the register \mathbf{r} .

base + displacement, $\text{BDISP}(\cdot, \cdot) \mathbf{r} \mathbf{d}$. The first argument is a register \mathbf{r} containing a base address, the second argument is the displacement \mathbf{d} .

base + displacement deferred, $\text{BDISPDEF}(\cdot, \cdot) \mathbf{r} \mathbf{d}$. The address of the value accessed is computed in the same way as the value in base + displacement mode.

predecrement, $\text{PREDEC}(\cdot, \cdot) \mathbf{r}$. Similar to register deferred mode, except that the address in the register \mathbf{r} is decreased before the value is accessed.

postincrement, $\text{POSTINC}(\cdot, \cdot) \mathbf{r}$. Again similar to register deferred. The address in \mathbf{r} is increased after the value is accessed.

immediate, $\text{IMM} \mathbf{v}$. The argument \mathbf{v} is the value accessed.

$$\begin{aligned}
& \{1.5\} \\
& = \text{arith}(x, y, r) \text{ rop}(x) \text{ wop}(a, y) \\
& \quad \quad \quad \Rightarrow \quad \quad \quad a[r]. \\
& \{1.4 \text{ arith}(x, y, y - x) = \text{SUB}. \} \\
& = \text{SUB} \text{ rop}(x) \text{ wop}(a, y) \quad \quad \quad r = y - x \\
& \quad \quad \quad \Rightarrow \quad \quad \quad a[r]. \\
& \{1.2 \text{ rop}(v_1) = \text{POSTINC}(-, v_1). \} \\
& = \text{SUB} \text{ POSTINC}(-, v_1) \text{ wop}(a, y) \quad \quad \quad x = v_1 \\
& \quad \quad \quad \Rightarrow \quad \quad \quad a[r]. \\
& \{1.1 \text{ POSTINC}(b_1, v_2) \text{ r}_1 = \text{r}_1[b_1] \text{ b}_1[v_2] \text{ c}_1 \Rightarrow \text{r}_1[\text{c}_1]. \} \\
& = \text{SUB} (\text{POSTINC} \text{ r}_1) \text{ wop}(a, y) \quad \quad \quad - = b_1, v_1 = v_2 \\
& \quad \quad \quad \text{r}_1[b_1] \text{ b}_1[v_2] \text{ c}_1 \\
& \quad \quad \quad \Rightarrow \text{r}_1[\text{c}_1] \quad \quad \quad a[r]. \\
& \{1.2 \text{ wop}(a_3, v_3) = \text{REGDEF}(a_3, v_3). \} \\
& = \text{SUB} (\text{POSTINC} \text{ r}_1) \text{ REGDEF}(a_3, v_3) \quad \quad \quad a = a_3, y = v_3 \\
& \quad \quad \quad \text{r}_1[b_1] \text{ b}_1[v_2] \text{ c}_1 \\
& \quad \quad \quad \Rightarrow \text{r}_1[\text{c}_1] \quad \quad \quad a[r]. \\
& \{1.1 \text{ REGDEF}(b_4, v_4) \text{ r}_2 = \text{r}_2[b_4] \text{ b}_4[v_4]. \} \\
& = \text{SUB} (\text{POSTINC} \text{ r}_1) (\text{REGDEF} \text{ r}_2) \quad \quad \quad a_3 = b_4, v_3 = v_4 \\
& \quad \quad \quad \text{r}_1[b_1] \text{ b}_1[v_2] \text{ c}_1 \text{ r}_2[b_4] \text{ b}_4[v_4] \\
& \quad \quad \quad \Rightarrow \text{r}_1[\text{c}_1] \quad \quad \quad a[r].
\end{aligned}$$

Figure 1: Archetype expansion in the linear machine.

resolution set) are resolved. Substitutions for variables in the textual expansion, defined by these resolved equations, then result in the update scheme above.

Textual expansion. The arithmetic command scheme 1.5 expands as shown in figure 1. The archetypes, and their index numbers are included in figure 1 for clarity. The equations at the right are the equations added at each step to the resolution set, which is initially empty. The final expansion is

$$\begin{aligned}
& [1.12] \text{ SUB} (\text{POSTINC} \text{ r}_1) (\text{REGDEF} \text{ r}_2) \\
& \quad \quad \quad \text{r}_1[b_1] \text{ b}_1[v_2] \text{ c}_1 \text{ r}_2[b_4] \text{ b}_4[v_4] \\
& \quad \quad \quad \Rightarrow \text{r}_1[\text{c}_1] \quad \quad \quad a[r].
\end{aligned}$$

and a substitution must be found for the non-ground variables \mathbf{r} and \mathbf{a} .

Resolution. The resolution set is $\{r = x - y, x = v_1, - = b_1, v_1 = v_2, a = a_3, y = v_3, a_3 = b_4, v_3 = v_4\}$. The expressions in the resolution set can then be resolved as shown in figure 2, to give substitutions for \mathbf{r} and \mathbf{a} . In this schematic representation of resolution, arrows (\rightarrow) represent grounding due to expressions in the update scheme, and equality ($=$) grounding via an equation from the resolution set. For example, if r_1 is ground, then b_1 is also ground, due to the locator expression $\mathbf{r}_1[b_1]$, and if v_4 is ground then v_3 is ground, due to the grounding set equation $v_3 = v_4$. An asterisk (*) indicates an equality that will be used in the substitution. Note that,

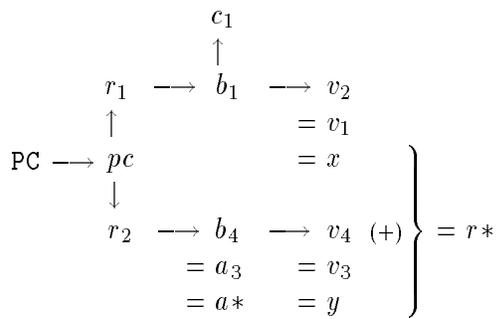
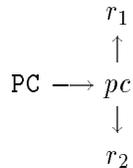


Figure 2: Resolving expressions in the linear machine.

for the sake of brevity, the initial part of the diagram in figure 2,



could have been omitted, as a similar part of diagram 8 on page VI.6 in appendix VI is.

Substitution. Substituting the values obtained in 1.12 gives

$$\begin{array}{l}
 \text{SUB (POSTINC } r_1) \text{ (REGDEF } r_2) \\
 r_1[b_1] \ b_1[v_2]c_1 \ r_2[b_4] \ b_4[v_4] \\
 \implies r_1[c_1] \ \quad \quad \quad b_4[v_4 - v_2].
 \end{array}$$

which is obviously equivalent to scheme 1.11.

2 The Tree Machine

The machine in section 1 will now be extended. The result is a machine for a tree structured language. In this language operands have a tree like structure, in which subtrees represent intermediate results. The update schemes specifying this machine are identical to those of the linear machine. The specifications differ only in their archetype definitions. For example, the update scheme

$$[1.3] \text{ MOV rop}(x) \text{ wop}(a, _)\implies a[x].$$

taken from the linear machine, is also part of the specification of the tree machine. The archetype definitions permit the representation of intermediate results. One possible expansion of scheme 1.3 is

$$\begin{array}{l}
 [2.1] \text{ MOV (IBDISP (IMOV (BDISP } r \ d_1) \text{ (IREG NIL)) } d_2) \\
 \text{ (IREGDEF (IMOV (IMM } a) \text{ (IREG NIL)))} \\
 r[b_1] \ b_1+d_1[b_2] \ b_2+d_2[v] \implies a[v].
 \end{array}$$

(where brackets have been added to improve legibility). The constants IBDISP and IREGDEF are the “intermediate” versions of BDISP and REGDEF, respectively, taking as their base value the result of an intermediate

“command” such as `IMOV`. For example, the value of the first operand is fixed by first determining the value accessed by the innermost addressing mode `BDISP r d1`. In specification 2.1 this is the value of `b2`. The archetype definitions do not require a value to be accessed by `IREG NIL`. The value `b2` is then passed by the `IMOV` command to the outermost addressing mode (`IBDISP (...) d2`), there to serve as base value.

Details of the expansion are given in figures 7 to 9 on pages VI.5 to VI.6 in appendix VI. The specification is given below, and summarised in figures 1 and 3 of appendix VI.

Although the full specification of the tree machine includes all of the update schemes and many of the archetypes from the specification of the linear machine they are repeated here for the sake of totality. Archetypes and update schemes unchanged from section 1 can be identified by their retention of the numbering given to them in that section. Constants which occur in both the linear machine and the tree machine, e.g. `ADD` and `BDISP`, are referred to as *linear constants*, and all others, e.g. `IMOV` and `IREG`, as *non-linear constants*. An expansion of a tree machine archetype is said to be *linear* if it contains no non-linear constants, and *non-linear* otherwise.

Addressing Modes. Intermediate addressing modes are added to the addressing mode archetypes. As in 1.1 the locator `b` in all of these archetype definitions is the base value of the addressing mode.

[1.1] `REG(r, b) r` = `r[b]`.
 `REGDEF(b, v) r` = `r[b] b[v]`.
 `BDISP(b+d, v) r d` = `r[b] b+d[v]`.
 `BDISPDEF(a, v) r d` = `r[b] b+d[a] a[v]`.
 `PREDEC(a, v) r` = `r[b] a[v]b` \implies `r[a]`.
 `POSTINC(b, v) r` = `r[b] b[v]c` \implies `r[c]`.

[2.2] `IREG(_, b) ir(b)` = `.`.
 `IREGDEF(b, v) ir(b)` = `b[v]`.
 `IBDISP(b+d, v) ir(b) d` = `b+d[v]`.
 `IBDISPDEF(a, v) ir(b) d` = `b+d[a] a[v]`.

The archetype `ir(·)` used here is defined on page 78.

The classes of addressing modes in 1.2 must now be redefined, and a new class, `tmp(·)`, added. In fact two new classes are added, but the class `kern(·, ·)` is merely added for convenience, making the definitions of the other classes shorter.

[2.3] $\text{kern}(a, v) = \text{REGDEF}(a, v). \quad \text{rop}(v) = \text{kern}(_, v).$
 $\quad = \text{BDISP}(a, v). \quad = \text{IMM } v.$
 $\quad = \text{BDISPDEF}(a, v). \quad = \text{REG}(_, v).$
 $\quad = \text{IBDISP}(a, v). \quad = \text{POSTINC}(_, v).$
 $\quad = \text{IBDISPDEF}(a, v). \quad = \text{IREG}(_, v).$
 $\quad = \text{IREGDEF}(_, v).$

$\text{wop}(a, v) = \text{kern}(a, v). \quad \text{adr}(a) = \text{kern}(a, _).$
 $\quad = \text{REG}(a, v). \quad \text{trg}(t) = \text{LAB } t.$
 $\quad = \text{PREDEC}(a, v). \quad \text{tmp}(t) = \text{IREG}(_, t).$
 $\quad = \text{IREGDEF}(a, v).$

The $\text{trg}(\cdot)$ class, which appears as the operand of the “jump” commands is restricted to include only (symbolic) labels. This will simplify the proof in section 4.2. A suggestion for relaxing this restriction is given in section 4.3.

Intermediate Results. The machine is extended with “intermediate” versions of the arithmetic and data transfer commands.

[2.4] $\text{iarith}(y+x) = \text{IADD rop}(x) \text{ tmp}(y).$
 $\quad \text{iarith}(y-x) = \text{ISUB rop}(x) \text{ tmp}(y).$
 $\quad \text{iarith}(y * x) = \text{IMUL rop}(x) \text{ tmp}(y).$
 $\quad \text{iarith}(y/x) = \text{IDIV rop}(x) \text{ tmp}(y) = [x \neq 0] \Rightarrow .$

[2.5] $\text{imov}(x) = \text{IMOV rop}(x) (\text{IREG NIL}).$
 $\quad \text{imov}(a) = \text{IMEA adr}(a) (\text{IREG NIL}).$

The second operand of IMOV and IMEA is actually superfluous, but makes the definition, given in section 3, of the transformation of tree code to linear code simpler. The addressing mode IREG NIL is *not* an expansion of the $\text{IREG}(\cdot, \cdot)$ archetype, and does not occur anywhere other than here. Its rôle in tree code is as a “dummy”, a placeholder for the temporary register that may be assigned for an intermediate result in the corresponding linear code.

The intermediate arithmetic and data transfer commands can now be combined to form the class of intermediate results, which was used in the definition of intermediate addressing modes in 2.2.

[2.6] $\text{ir}(r) = \text{iarith}(r).$
 $\quad = \text{imov}(r).$

The remainder of the archetype definitions and update schemes are inherited from the linear machine.

Data Transfer Commands.

[1.3] $\text{MOV rop}(x) \text{ wop}(a, _) \Rightarrow a[x].$
 $\quad \text{MEA adr}(x) \text{ wop}(a, _) \Rightarrow a[x].$

Arithmetic Commands.

$$\begin{aligned}
\langle \text{tree} \rangle &\rightarrow \langle \text{dyadic} \rangle \langle \text{op} \rangle \langle \text{op} \rangle | \\
&\quad \langle \text{monadic} \rangle \langle \text{op} \rangle | \\
&\quad \langle \text{niladic} \rangle | \\
&\quad \text{NIL} \\
\langle \text{dyadic} \rangle &\rightarrow \langle \text{arith} \rangle | \langle \text{iarith} \rangle | \\
&\quad \langle \text{move} \rangle | \langle \text{imove} \rangle | \\
&\quad \langle \text{bool} \rangle \\
\langle \text{monadic} \rangle &\rightarrow \langle \text{jump} \rangle \\
\langle \text{niladic} \rangle &\rightarrow \text{RET} \\
\langle \text{arith} \rangle &\rightarrow \text{ADD SUB MUL DIV} \\
\langle \text{iarith} \rangle &\rightarrow \text{IADD ISUB IMUL IDIV} \\
\langle \text{move} \rangle &\rightarrow \text{MOV MEA} \\
\langle \text{imove} \rangle &\rightarrow \text{IMOV IMEA} \\
\langle \text{bool} \rangle &\rightarrow \text{CGT CGEQ CEQ CNE CLEQ CLT} \\
\langle \text{jump} \rangle &\rightarrow \text{JT} | \text{JF} | \text{JMP} | \text{JSR} \\
\langle \text{op} \rangle &\rightarrow \langle \text{dop} \rangle | \langle \text{iop} \rangle \\
\langle \text{dop} \rangle &\rightarrow \langle \text{dmode}_1 \rangle \langle x \rangle | \langle \text{dmode}_2 \rangle \langle r \rangle \langle d \rangle \\
\langle \text{dmode}_1 \rangle &\rightarrow \text{REG REGDEF PREDEC POSTINC LAB IMM} \\
\langle \text{dmode}_2 \rangle &\rightarrow \text{BDISP BDISPDEF} \\
\langle \text{iop} \rangle &\rightarrow \langle \text{imode}_1 \rangle \langle \text{tree} \rangle | \langle \text{imode}_2 \rangle \langle \text{tree} \rangle \langle d \rangle \\
\langle \text{imode}_1 \rangle &\rightarrow \text{IREG IREGDEF} \\
\langle \text{imode}_2 \rangle &\rightarrow \text{IBDISP IBDISPDEF}
\end{aligned}$$

Figure 3: A grammar for tree code.

terminals $\langle x \rangle$, $\langle r \rangle$ and $\langle d \rangle$ are not specified in this grammar. Though the definition of $\mathcal{T}[\cdot]$ is based on the grammar in figure 3, the update schemes and archetype definitions may impose further restrictions on configurations, and this will be indicated as necessary. Reference will also be made to the archetype definitions and update schemes where this is necessary to restrict the domain under consideration.

In the following definitions the prime symbol $'$ is the postfix operator which translates all tree machine constants to their linear machine counterparts, i.e. $\text{MOV}' = \text{MOV}$, $\text{IADD}' = \text{ADD}$, $\text{IMM}' = \text{IMM}$, $\text{IBDISP}' = \text{BDISP}$, etc.

$\mathcal{T}[\cdot]$ is defined in terms of $\mathcal{IR}[\cdot]$ (defined on page 81), and $\mathcal{M}[\cdot]$ (on

$$\begin{aligned}
T[\textit{dyadic } op_1 \textit{ } op_2] n &= \mathcal{IR}[op_2] n \\
&\quad \mathcal{IR}[op_1] n + 1 \\
&\quad \textit{dyadic}' (\mathcal{M}[op_1] n + 1) (\mathcal{M}[op_2] n) \\
T[\textit{monadic } op] n &= \mathcal{IR}[op] n \\
&\quad \textit{monadic}' (\mathcal{M}[op] n) \\
T[\textit{niladic}] n &= \textit{niladic}' \\
T[\textbf{NIL}] n &= \epsilon
\end{aligned}$$

Of the auxiliary functions $\mathcal{M}[\cdot]$ is the simpler. It gives the linear machine addressing mode corresponding to its (tree machine) argument.

$$\begin{aligned}
\mathcal{M}[\textit{dmode}_1 x] n &= \textit{dmode}_1 x \\
\mathcal{M}[\textit{dmode}_2 r d] n &= \textit{dmode}_2 r d \\
\mathcal{M}[\textit{imode}_1 ir] n &= \textit{imode}'_1 n \\
\mathcal{M}[\textit{imode}_2 ir d] n &= \textit{imode}'_2 n d
\end{aligned}$$

The function $\mathcal{IR}[\cdot]$ distinguishes between linear and non-linear addressing modes, returning code to compute an intermediate result, if necessary, by calling $T[\cdot]$.

$$\begin{aligned}
\mathcal{IR}[\textit{dmode}_1 x] n &= \epsilon \\
\mathcal{IR}[\textit{dmode}_2 r d] n &= \epsilon \\
\mathcal{IR}[\textit{imode}_1 ir] n &= T[ir] n \\
\mathcal{IR}[\textit{imode}_2 ir d] n &= T[ir] n
\end{aligned}$$

Note that $\mathcal{M}[op] n$ and $\mathcal{IR}[op] n$ are complementary in that $\mathcal{IR}[\cdot]$ produces code which “computes”, if necessary, the base value of op and places it in a register for the use of $\mathcal{M}[op]$, which applies the required indirections and displacements to provide the effective address (if any) and value of op .

The main transformation function $T[\cdot]$ can be seen as a specification of a (simple) register allocation algorithm for a machine with an unlimited supply of registers. The brackets included, to increase legibility, in the **MOV** instruction on page 76 are essential for correct parsing of the argument of $T[\cdot]$. Correct bracketing will not be explicitly detailed here, but will be considered to be obvious.

4 Proving Semantic Equivalence

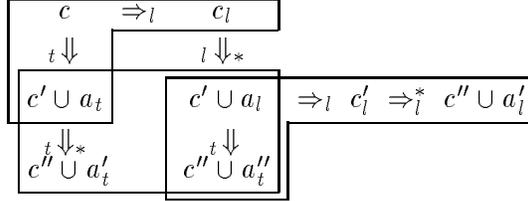
In addition to $\mathcal{IR}[\cdot]$ it is assumed that, for any given configuration, there is a function from locators in the tree code programme store to locators in the linear code programme store, such that symbolic labels are preserved. I.e. if $\mathcal{L}[\cdot]$ is this function, and $\mathbf{l}[\mathbf{c}]\mathbf{r}$ is part of the (tree code) programme store, then

$$\mathcal{L}[\mathbf{l}][T[\mathbf{c}]]\mathcal{L}[\mathbf{r}]$$

will be part of the (linear code) programme store derived by applying $T[\cdot]$ and $\mathcal{L}[\cdot]$. Note that $T[\cdot]$ is defined for individual tree machine commands. When applying it to a tree code programme store it will be necessary to map it across the sequence of instructions in the programme store.

It is claimed that $\mathcal{L}[\cdot]$ and $\mathcal{T}[\cdot]$, applied to the programme store and the contents of PC, when combined with the identity function for all other components of the configuration, form a correct translation, partial outside **Aux**, from configurations of the tree machines to configurations of the linear machine.

The proof is as sketched in chapter six. It will be shown that $\mathcal{T}[\cdot]$ defines a local translation, partial outside **Aux**. Since **Aux** is trivially irrelevant, both locally and transitively, to the tree machine (it is a no-access memory), transitive translation can be established by inductive diagram chasing similar to that in section 5.1 of chapter six:



where \Rightarrow_t is an update in the tree machine, and \Rightarrow_l an update in the linear machine, and a_t, a_l, a'_t, a'_l and a''_t are all subsets of **Aux**.

The proof will establish that $(c \Rightarrow_t c' \cup a_t) \implies (c \Rightarrow_l^* c' \cup a_l)$. The remaining non-trivial part of the translation box diagram

$$(c \Rightarrow_l c_l) \implies (\exists c' : c_l \Rightarrow_l^* c' \cup a_l \wedge c \Rightarrow_t c' \cup a_t)$$

follows automatically from the deterministic nature of the linear machine.

At certain stages in the proof it will be necessary to consider translation partial outside a subset of **Aux**. The subset will always consist of those cells in **Aux** to the right of the locator $\mathbf{AUX} + n$, for some n . Such a partial translation will be indicated by subscripting the translation symbol with n — i.e. $\triangleright_n \equiv \triangleright_{\{\mathbf{AUX}+k | k > n\}}$.

The nonterminals of the grammar in figure 3 will be used consistently throughout the proof — a variable ‘*iop*’ represents an expansion of the nonterminal $\langle iop \rangle$. Where necessary subscripts will be used to distinguish objects of the same type, e.g. op_1 and op_2 .

4.1 An Important Lemma

The key to the proof is the following lemma:

Lemma For any non-linear operand op

$$\mathcal{IR}[op] n \triangleright_n \text{MOV}(\text{IMM } b) (\text{REG } n)$$

where b is the base value of op . \diamond

If op is non-linear, i.e. an *iop*, it must, in accordance with the grammar, be one of

IREG *tree*
 IREGDEF *tree*
 IBDISP *tree d*
 IBDISPDEF *tree d*

where, from archetype definitions 2.2, 2.6, 2.4 and 2.5, and 2.3, *tree* is one of

```

iarith dop1 iop2
iarith iop1 iop2
imove dop1 IREG NIL
imove iop1 IREG NIL.

```

The occurrences of *iop₁* and *iop₂* call for a proof by induction. The proof will only be given for the second expansion of *tree* (*iarith iop₁ iop₂*), the other choices — including the third, which is the induction base — being analogous. Even more specifically, of all 48 possible choices for *op* only

```
IREGDEF (IADD iop1 iop2)
```

will be considered.

The proof is simple. Firstly,

$$\begin{aligned}
\mathcal{IR}[op] n &= \mathcal{IR}[\text{IREGDEF (IADD } iop_1 iop_2)] n \\
&= \mathcal{T}[\text{IADD } iop_1 iop_2] n \\
&= \mathcal{IR}[iop_2] n \\
&\quad \mathcal{IR}[iop_1] n + 1 \\
&\quad \text{ADD } (\mathcal{M}[iop_1] n + 1) (\mathcal{M}[iop_2] n).
\end{aligned}$$

Assuming the base values of *iop₁* and *iop₂* to be *b₁* and *b₂* respectively gives, applying the induction hypothesis,

$$\mathcal{IR}[iop_2] n \triangleright_n \text{MOV (IMM } b_2) (\text{REG } n)$$

and

$$\mathcal{IR}[iop_1] n + 1 \triangleright_{n+1} \text{MOV (IMM } b_1) (\text{REG } n + 1)$$

The only cell covered by the \triangleright_{n+1} relation and not by the \triangleright_n relation is register *n + 1*. The contents of this cell will be updated by the **MOV (IMM *b₁*) (REG *n + 1*)** command (see below) and therefore

$$\begin{aligned}
\mathcal{IR}[op] n \triangleright_{n+1} &\text{MOV (IMM } b_2) (\text{REG } n) \\
&\text{MOV (IMM } b_1) (\text{REG } n + 1) \\
&\text{ADD } (\mathcal{M}[iop_1] n + 1) (\mathcal{M}[iop_2] n).
\end{aligned}$$

It follows, from archetype definitions 2.4 and 2.3, that *iop₂* can only be **IREG *tree₂***, while, by the grammar, there are again four possible choices for *iop₁*. Again only one case will be considered, **IBDISP *tree₁* *d₁***, the other cases again being analogous. Then $\mathcal{M}[iop_1] n + 1$ is **BDISP *n + 1 d₁***, and $\mathcal{M}[iop_2] n$ is **REG *n***.

Now *iop₁* and *iop₂* are expansions of archetypes **IBDISP(*b₁ + d₁*, *r₁*)** and **REG(*-, b₂*)**, and any update scheme covering *op* will contain, inter alia, the contexts of these (expanded) archetypes, and in particular, on the left hand side, the locator expression *b₁ + d₁[v₁]*.

Now any configuration satisfying

$$\begin{array}{l}
\text{PC[pc]} \quad \text{pc}[\text{MOV (IMM } b_2) (\text{REG } n)] \sim \\
\quad \text{qc}[\text{MOV (IMM } b_1) (\text{REG } n + 1)] \sim \\
\quad \text{rc}[\text{ADD (BDISP } n + 1 d_1) (\text{REG } n)] \text{sc} \quad b_1 + d_1[v_1]
\end{array}$$

updates deterministically to one satisfying

$$\text{PC}[\text{rc}] \quad \text{rc}[\text{ADD} (\text{BDISP } n + 1 \ d_1) (\text{REG } n)]\text{sc}$$

$$b_1 + d_1[v_1] \quad n + 1[b_1] \quad n[b_2]$$

and then to one satisfying

$$\text{PC}[\text{sc}] \quad n[v_1 + b_2].$$

Clearly, then,

$$\mathcal{IR}[\text{op}] \ n \triangleright_n \text{MOV} (\text{IMM } v_1 + b_2) (\text{REG } n).$$

Furthermore, according to the definition of the tree machine, $v_1 + b_2$ is indeed the base value of

$$\text{IREGDEF} (\text{IADD} (\text{IBDISP } tree_1 \ d_1) (\text{IREG } tree_2))$$

where $tree_1$ and $tree_2$ are expansions of $\text{ir}(b_1)$ and $\text{ir}(b_2)$ respectively. This establishes the lemma.

4.2 The Proof

Tree code is either of the form *dyadic op op*, or of the form *monadic op*, or of the form *niladic*. Examination of the update schemes suffices to show that $\mathcal{T}[\text{niladic}] \ n \triangleright_n \text{niladic}$. The operand of a *monadic op* is restricted by the archetype definitions to a **LAB** \mathfrak{t} . This case reduces to the case for *niladic*, as does a *dyadic dop dop*. The most general case is *dyadic iop iop*. The proof that $\mathcal{T}[\text{dyadic iop iop}] \ n \triangleright_n \text{dyadic iop iop}$ is very similar to the proof of the lemma in section 4.1, though it no longer requires induction, and will not be presented. The remaining two cases — i.e. a *dyadic op op* in which one of the operands is a *dop* — are, in structure at least, simpler. Once it has been established that $\mathcal{T}[\text{niladic}] \ n \triangleright_n \text{niladic}$, that $\mathcal{T}[\text{monadic op}] \ n \triangleright_n \text{monadic op}$ and that $\mathcal{T}[\text{dyadic op op}] \ n \triangleright_n \text{dyadic op op}$ it is trivial to show that $\mathcal{T}[\text{code}] \ 0 \triangleright^{\text{aux}} \text{code}$. Care should be taken with the programme flow commands, *jump*. The archetypes require the operand of a *jump* to be an expansion of $\text{trg}(\mathfrak{t})$. In the tree machine this must be a label — a locator in the programme store of the tree machine. Preservation of labels then ensures the required semantic equivalence. It is for this reason that $\text{trg}(\mathfrak{t})$ is restricted to expand to **LAB** \mathfrak{t} in the tree machine.

4.3 Possible Extensions

Stacks. The transformation given assumes a sufficient supply of auxiliary registers. An alternative definition could provide translations in which intermediate results are passed via a stack, if there are insufficient auxiliary registers. The transformation to code for passing results via the stack is simply presented here, without a proof of its correctness. This is, of course, not the only possible definition, nor even probably the most “preferable”. It is however almost certainly the simplest. The proof is roughly along the same lines as that already presented.

$$\begin{aligned} \mathcal{T}[\text{dyad } op_1 \ op_2] \ n &= \mathcal{IR}[op_1] \ n \\ &\quad \text{MOV} (\mathcal{M}[op_1] \ n) (\text{PREDEC SP}) \\ &\quad \mathcal{IR}[op_2] \ n \\ &\quad \text{dyad}' (\text{POSTINC SP}) (\mathcal{M}[op_2] \ n) \end{aligned}$$

This transformation can either be added to the definition of $\mathcal{T}[\cdot]$, to be called whenever the number of auxiliary registers available falls below a critical reserve, or used to make $\mathcal{T}[\cdot]$ a relation specifying both code for passing intermediate results via auxiliary registers and code in which results are passed via the stack. In both cases, proving that the new transformation preserves the semantics suffices, in combination with the proof above, to ensure the correctness of the translation (or translations) provided by $\mathcal{T}[\cdot]$.

Another possibility is to have $\mathcal{T}[\cdot]$ return some error value when the auxiliary registers and/or the stack are exhausted. Correct translation must then be proved for any code for which $\mathcal{T}[\cdot]$ does not return this error value.

Jumps. Restricting the targets of jumps, in particular subroutine jumps, to symbolic labels is fairly drastic, making it impossible, for example, to return procedures or functions as results — or, at least, to call a procedure via such a result. This restriction can be relaxed fairly simply.

Both the linear machine and the tree machine have a programme store, say `ProgStore`. Since `ProgStore` is a programme store any `t` in a `trg(t)` will be of type `ProgStore`. The property required of the target of a jump is not that it *be* a (symbolic) label, but that it *have the value* of such a label. Restricting the target of a jump to a `LAB t` corresponds to requiring that `t` be a label. The lesser restriction can be achieved by requiring all objects of type `ProgStore` in the initial configuration to be symbolic labels, and, furthermore, requiring that no object of type `ProgStore` in the update plan be any expression more complicated than a constant (a symbolic label), or a single variable. Implicitly, such objects may therefore not appear as an argument of, for example, an `ADD` or `IADD`. This ensures that labels may be copied, and moved from cell to cell, but that no “new” `ProgStore` objects can be created. In any update script satisfying these conditions almost any addressing mode may be used in the operand of a jump instruction.

5 Conclusions

Update Plans, and in particular the archetype mechanism, make it possible to reason formally about low level activities. Transformations at this level can be proven correct. While the proof outlined in this chapter is fairly baroque, it is hard to see how any possible proof could be simpler. The combinatorial explosion of addressing modes almost certainly makes extensive case analysis unavoidable. Update Plans, and in particular the archetype mechanism, at least offer the means to emphasise the similarities, rather than the differences, between the cases, and possibly open the way to (partial) automation of the proof obligation.

Register allocation is but one transformation. The methods presented here can also be applied to other such transformations, be it at tree code level (with a slight abuse of notation):

$$dyad (IREG (IMOV op_1(IREGNIL))) op_2 \equiv dyad op_1 op_2,$$

or at linear machine level where a peephole optimiser could be defined.

Since all these applications preserve semantics they can also be combined; first optimising at the tree-code level, then performing register allocation, and then optimising the resultant linear code.

While nondeterminism has been mentioned in passing in chapters one and seven, and specified implicitly in chapter three, it has not yet been covered in detail. Implementational aspects of nondeterminism are covered in chapter ten. This chapter is about the relation of nondeterminism to synchronous parallelism.

In chapter seven nondeterminism in Update Plans was exploited for the specification of asynchronous parallel processors. An adaptation of nondeterminism can be used to specify synchronous parallelism. This adaptation is described in sections 1, 2 and 3. An application — pipelining in a partial specification of the Berkeley RISC II CPU — is presented in section 4.

1 Informal Introduction

Update Plans as defined in chapters two and three, and as used in chapters four and eight, intuitively “work” by instantiating all applicable update schemes and then making a nondeterministic choice which of the update rules thus obtained to apply. An alternative would be to apply all of the update rules simultaneously. This is the idea behind a *parallel block*. A parallel block is a set of update schemes all applicable instantiations of which will be applied at the same time, if possible. The caveat is that some of these instantiations may have conflicting right hand sides. If this is the case, then none of the schemes are applied, the reasoning being that the update rules are applied as if they all formed one update rule, and update rules with inconsistent right hand sides may not be applied, as detailed in chapter three.

2 Syntax

Parallel blocks are delimited by the *open parallel block symbol*, ‘(|’, and the *close parallel block symbol*, ‘|)’. The pipeline symbol ‘|’ may be used anywhere in an update plan that whitespace is permitted. This can be used to emphasise a parallel block by prefixing each line with a pipeline

symbol.

Example 1

The parallel block

$$\begin{aligned} & (| \\ & \quad \text{IR+S[src(x)]} \implies \text{RX[x]}. \\ & \quad \text{IR+D[dst(ea, y)]} \implies \text{RY[y] MAR[ea]}. \\ & |) \end{aligned}$$

may be written

$$\begin{aligned} & (| \\ & | \quad \text{IR+S[src(x)]} \implies \text{RX[x]}. \\ & | \quad \text{IR+D[dst(ea, y)]} \implies \text{RY[y] MAR[ea]}. \\ & |) \end{aligned}$$

or, making use of typesetting possibilities

$$\left(\begin{array}{l} | \\ | \quad \text{IR+S[src(x)]} \implies \text{RX[x]}. \\ | \quad \text{IR+D[dst(ea, y)]} \implies \text{RY[y] MAR[ea]}. \\ | \end{array} \right)$$

In the last case, and under the condition that this is the only use of the pipeline symbol, the open and close parallel block symbols may be omitted, giving

$$\begin{array}{l} | \quad \text{IR+S[src(x)]} \implies \text{RX[x]}. \\ | \quad \text{IR+D[dst(ea, y)]} \implies \text{RY[y] MAR[ea]}. \end{array}$$

The grammar given in chapter two, and modified in chapters five and seven must again be extended. Add the production rules

$$\langle \text{item} \rangle \rightarrow \langle \text{parallel block} \rangle$$
$$\langle \text{parallel block} \rangle \rightarrow (| \langle \text{alternatives} \rangle |)$$

The full grammar of Update Plans, with all extensions, is given in appendix II.

3 Formal Specification

In chapter three the semantics of Update Plans was defined in terms of applications of update schemes: the update relation \Rightarrow_s is defined in terms of the function \mathcal{S} (which given a left hand side, a guard and a configuration gives a set of substitutions under which an update scheme having that left hand side and guard will be applicable to the given configuration), and the interpretation function $\mathcal{I}[\cdot]$ (which interprets the instantiation of a right hand side). In the definition of \Rightarrow_s on page 22 the relation can be said to

create and apply applicable update rules “on the fly”. An equivalent view is to first derive applicable update rules and to define the semantics of Update Plans in terms of update rule applications.

Given an update plan, P , and a configuration, c , the set of update rules from P applicable to c is given by

$$\mathcal{R} P c = \{(l^\sigma, r^\sigma) \mid \exists g : (l, g, r) \in P, \sigma \in \mathcal{S} \text{ l g c and } r^\sigma \text{ is consistent}\}.$$

This can be simplified to give the set of updates (right hand sides), since the left hand sides are only relevant for determining which instantiations of which update schemes are applicable.

$$\mathcal{U} P c = \{r \mid (l, r) \in \mathcal{R} P c\}.$$

The relation between configurations defined by an update plan P can now be expressed as

$$c \Rightarrow_P c' \equiv \exists r \in \mathcal{U} P c : r \text{ is consistent} \wedge c' = \mathcal{I}[[r]] \uplus c.$$

This is equivalent to the definition of \Rightarrow_P given in chapter three but gives, perhaps, a clearer picture of the possibilities from which a nondeterministic choice is made during evaluation of an update script.

The semantics of parallel blocks can be defined in the same way. Rather than making a nondeterministic choice between applicable update rules all such rules are applied simultaneously. In terms of \mathcal{U} , an update by parallel block B is defined by

$$c \Rightarrow_B c' \equiv \bigcup (\mathcal{U} B c) \text{ is consistent} \wedge c' = \mathcal{I} \left[\left[\bigcup (\mathcal{U} B c) \right] \right] \uplus c.$$

For the case that $\mathcal{U} B c$ is a singleton set this reduces to the definition of \Rightarrow_s given in chapter three.

Example 2

If the set of applicable update rules from a parallel block consists of the update rules

$$\text{IR+S [REG RO] RO [X] } \Longrightarrow \text{RX [X] .}$$

and

$$\text{IR+D [REG R1] R1 [Y] } \Longrightarrow \text{RY [Y] MAR [R1] .}$$

these will be applied simultaneously as if the update rule

$$\begin{aligned} &\text{IR+S [REG RO] RO [X] IR+D [REG R1] R1 [Y]} \\ &\Longrightarrow \text{RX [X] RY [Y] MAR [RY] .} \end{aligned}$$

were to be applied.

In keeping with the macro character of archetypes, archetype expansion conceptually takes place before interpretation of parallel blocks, and any expansions remain within the parallel block.

Short-Immediate Format

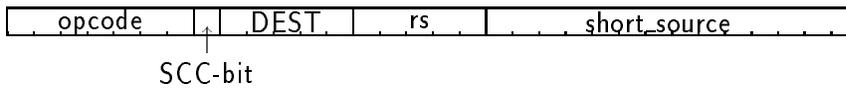
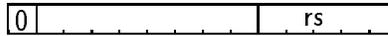


Figure 1: Format of a RISC II arithmetic instruction.

Register Source



Immediate Source



Figure 2: Short-source instruction field formats.

4 An Example: The Berkeley RISC II CPU

In the following example a subset of the Berkeley RISC II instruction set will be specified, first at the instruction level, to provide a general view; then at the instruction cycle level, as in chapter seven, but without pipelining; and then, finally, with pipelining, first without internal forwarding, and then with internal forwarding. Only a small subset of the instruction set is specified, in order to limit the length of this example. The specification of this subset will, however, be fairly detailed, and will also serve as an illustration of the use of Update Plans for the specification of concrete machines. In particular, the instruction level specification is fairly long. It is, however, complete, for this subset. The specification is based on the description of the Berkeley RISC II CPU given by Katavenis [35].

4.1 Instruction Set

The subset that will be specified is that of the arithmetic commands.

An arithmetic instruction on the RISC II has a *short-immediate format*, as shown in figure 1 in which: **DEST** is the destination, which must be a register; **rs**, which must also be a register, is the first operand; and **short_source** is the second. The second operand, **short_source**, is a *short-immediate operand* which may have one of two formats. These formats are shown in figure 2.

Addressing Modes. The RISC II has three types of register — global registers shared by all routines, local registers addressable in a local window indicated by the current window pointer **CWP**, and register 0, which always has value 0. In the following definition register 0 will be considered a global register. It will be treated as a special case as the need arises.

$$\begin{aligned}\text{reg}(\text{GBASE}+\text{r}) &= \text{global}(\text{r}). \\ \text{reg}(\text{cwp}+\text{r}) &= \text{local}(\text{r}) \text{CWP}[\text{cwp}].\end{aligned}$$

GBASE is the left locator of register 0. The two archetypes `global(r)` and `local(r)` distinguish between global and local registers. The global registers are registers 0 to 9, while registers 10 to 31 are local. The archetypes are defined as:

$$\begin{aligned} \text{global}(r) &= r = [0 \leq r \leq 9] \Rightarrow . \\ \text{local}(r) &= r = [10 \leq r \leq 31] \Rightarrow . \end{aligned}$$

Note that the variable `r` in these archetype definitions will create a new variable in any update scheme in which the archetype occurs, said variable then obtaining its value by instantiation in the usual way.

A register source can now be defined.

$$\begin{aligned} \text{rs}(0) &= \text{reg}(\text{GBASE}). \\ \text{rs}(\text{val}) &= \text{reg}(\text{ea}) \text{ea}[\text{val}] = [\text{ea} \neq \text{GBASE}] \Rightarrow . \end{aligned}$$

The last instruction field of a short-immediate format instruction can have one of two formats, and this is specified by:

$$\begin{aligned} \text{short}(\text{val}) &= \text{REG } \text{rs}(\text{val}). \\ &= \text{IMM } \text{val}. \end{aligned}$$

Condition Codes. Condition codes are set if the SCC-bit is `0N`.

$$\begin{aligned} \text{scc}(_) &= \text{OFF}. \\ \text{scc}(\text{flags}) &= 0N \Rightarrow \text{CC}[\text{flags}]. \end{aligned}$$

For arithmetic operations the values of the condition codes, in the order `N`, `Z`, `V`, `C`, are given by

$$\text{aflg}(v) = (v <_s) (v = 0) \neg(\text{MIN}_s \leq_s v \leq_s \text{MAX}_s) (v >_u \text{MAX}_u) |.$$

The operators $<_s$, \leq_s (signed comparison operators), $>_u$ (unsigned comparison) and $=$ are assumed to be defined elsewhere, as are the constants `MAXs`, `MINs` and `MAXu`. In a complete, bit level, specification, the values of the flags might be determined with reference to the values of the source operand, as well as that of the result. Such a specification will not be given here.

The SCC-bit in an arithmetic instruction can now be represented by `scc(aflg(v))`, where `v` is the result of the arithmetic operation. This will expand to either

$$\text{OFF} \Rightarrow .$$

or

$$0N \Rightarrow \text{CC}[(v <_s) (v = 0) \neg(\text{MIN}_s \leq_s v \leq_s \text{MAX}_s) (v >_u \text{MAX}_u)].$$

Arithmetic Instructions. The arithmetic opcodes are defined analogously to those in chapter eight.

$$\begin{aligned} \text{arith}(x, y, x+y) &= \text{ADD}. \\ \text{arith}(x, y, x+y+c) &= \text{ADDC } C[c]. \\ \text{arith}(x, y, x-y) &= \text{SUB}. \\ \text{arith}(x, y, x-y-(1-c)) &= \text{SUBC } C[c]. \\ \text{arith}(x, y, y-x) &= \text{SUBI}. \\ \text{arith}(x, y, y-x-(1-c)) &= \text{SUBCI } C[c]. \end{aligned}$$

$$\begin{aligned} & \text{ADD OFF } r \text{ } s \text{ (IMM } y) \text{ CWP}[cwp] \text{ cwp}+s[x] \\ & = [r \in \text{GLBL} \wedge r \neq 0 \wedge s \in \text{LOC} \wedge cwp+s \neq \text{GBASE}] \Rightarrow \\ & \quad \text{GBASE}+r[y+x]. \end{aligned}$$

Figure 3: A typical Berkeley RISC II arithmetic instruction. This instruction takes the value y , adds it to the contents of local register s , and places the result in register r . The condition codes are not set.

The register C contains the carry bit. All arithmetic instructions are now defined by

$$\begin{aligned} & \text{arith}(x, y, r) \text{ scc}(\text{aflg}(r)) \text{ reg}(\text{ea}) \text{ rs}(x) \text{ short}(y) \\ & = [\text{ea} = \text{GBASE}] \Rightarrow . \\ & " \quad = [\text{ea} \neq \text{GBASE}] \Rightarrow \text{ea}[r]. \end{aligned}$$

A possible concrete example of this update scheme, with the archetypes fully expanded, is given in figure 3. Some expressions derived during the expansion have been simplified.

4.2 The Instruction Cycle

The instruction cycle consists of three phases; instruction fetch, execute, and write. In the instruction fetch phase the instruction currently addressed by the PC is copied to the instruction register IR. The execute phase accesses the operands and performs the operation, setting the condition codes, if necessary, and placing the result in the RES register. In order to keep all accesses of the internal structure of an instruction within one phase, the destination address is also copied to the DST register, where it will be used by the write phase. The execute phase also updates the PC to contain the address of the next instruction to be executed. The PC is only updated at execution, rather than when the instruction is fetched, since addressing may be relative to the PC for some instructions. Finally, the write phase copies the result from the RES register to the destination. As in chapter seven, a clock is introduced to ensure that the phases take place sequentially. In the next step, pipelining, the clock will be eliminated.

$$\begin{aligned} & \text{FETCH}() = \\ & \quad \text{PC}[pc] \text{ pc}[\text{instruction}] \Rightarrow \text{IR}[\text{instruction}]. \\ & \text{EXEC}() = \\ & \quad \text{IR}[\text{arith}(x, y, r) \text{ scc}(\text{aflg}(r)) \text{ reg}(\text{dst}) \text{ rs}(x) \text{ short}(y)] \text{ PC}[pc] \\ & \quad \Rightarrow \text{RES}[r] \text{ DST}[\text{dst}] \text{ PC}[pc+\text{WORD}]. \\ & \text{WRITE}() \\ & = \text{DST}[\text{dst}] \text{ RES}[r] = [\text{dst} = \text{GBASE}] \Rightarrow . \\ & = \quad " \quad = [\text{dst} \neq \text{GBASE}] \Rightarrow \text{dst}[r]. \end{aligned}$$

WORD is the length of an instruction. The instruction cycle is defined by:

$$\begin{aligned} & \text{tick}() = \text{FETCH}() \Rightarrow \text{EXEC}. \\ & = \text{EXEC}() \Rightarrow \text{WRITE}. \\ & = \text{WRITE}() \Rightarrow \text{FETCH}. \end{aligned}$$

while that of the pipelined specification would be still be

$$\text{PC}[\text{pc}_1] \Longrightarrow \text{PC}[\text{pc}_5] \text{ GBASE}+1[3] \text{ GBASE}+3[5].$$

This problem is avoided by using *internal forwarding*. The execution phase must check if the result waiting to be written should be used instead of a stale value. This can be done by revising the specification of a register source.

$$\begin{aligned} \text{rs}(0) &= \text{reg}(\text{GBASE}). \\ \text{rs}(\text{res}) &= \text{reg}(\text{dst}) \text{ DST}[\text{dst}] \text{ RES}[\text{res}] = [\text{dst} \neq \text{GBASE}] \Rightarrow . \\ \text{rs}(\text{val}) &= \text{reg}(\text{ea}) \text{ DST}[\text{dst}] \text{ ea}[\text{val}] = [\text{ea} \neq \text{GBASE} \wedge \text{ea} \neq \text{dst}] \Rightarrow . \end{aligned}$$

The definition of $\text{rs}(\cdot)$ has been extended. The second declaration “intercepts” a result of the previous instruction. DST and RES contain the result register and value of the previous instruction, which are still in the pipeline. If $\text{rs}(\cdot)$ requires the value in that result register it should take the “new” value from RES , rather than the “old” value at ea , since this will not yet have been updated.

The rest of the specification is unchanged. The techniques of chapter eight could be applied to these specifications in order to prove their semantic equivalence. Certain restrictions will have to be imposed on the code, in particular on instructions immediately after a jump instruction, in order to make this possible.

5 Conclusions

The combination of the archetype mechanism and a simple notation for indicating parallelism makes it possible to write short, concise specifications of pipelining architectures. Using methods similar to those applied in chapter eight it would be possible to derive clearly defined conditions for (generated) code which would guarantee equivalence of pipelined and non-pipelined code. It also becomes possible to derive provably correct optimisers which take advantage of pipelining.

The previous chapters of this thesis have defined Update Plans. In this chapter a possible implementation of the specification language will be presented. It is not the intention to provide a definitive specification of an implementation, but rather to illustrate techniques that could be used for such an implementation.

Section 1 introduces certain notational conventions specific to this chapter. Sections 2 and 3 discuss the implementation of deterministic Update Plans, section 2 for basic Update Plans as defined in chapters one, two and three, and section 3 for Update Plans with archetypes as defined in chapter seven. Section 4 discusses how backtracking could be implemented for nondeterministic Update Plans, and how this mechanism could be adapted to implement the parallel blocks from chapter nine. Finally, section 5 evaluates the implementation described, and suggests some techniques for improving its efficiency.

The key to the implementation is access to the (internal representations of) values of terms appearing in an update scheme. The implementation presented takes a straightforward approach, storing values, or pointers to values, in a *value store*. A more realistic approach would store some values in temporary registers. This is, however, irrelevant to the further implementation, as long as an access path to the values can be statically determined. The implementation first constructs a storage structure for values, then checks the guard for applicability and the right hand side for consistency and then, if these checks are successful, applies the update scheme.

1 Notes on Notation

This chapter applies certain notational conventions. These are summarised here.

Canonical Forms. The update plan to be implemented will be assumed to be in canonical form — i.e. all syntactic sugar will have been removed and all guards will be made explicit.

left locator $\mathbf{val} + \mathbf{0}_{obj}$ will contain (a pointer to) obj . In all situations, in the implementation presented here, in which such objects are addressed by way of their offset, the left locator of the storage structure will be in a register \mathbf{V} . An ambidextrous archetype can be defined to access these values.

$$\mathbf{val}(\mathbf{0}_{obj}) = \mathbf{val} \mid \mathbf{V}[\mathbf{v}] \mathbf{v} + \mathbf{0}_{obj}[\mathbf{val}].$$

The notation \mathbf{val}^{obj} will be used as an abbreviation for $\mathbf{val}(\mathbf{0}_{obj})$. This convention will also be applied to structures in other stores.

$|\mathbf{Locator}|$. The length of a locator occurs frequently in the update schemes in this chapter. ‘ \mathbb{L} ’ will be used as abbreviation for $|\mathbf{Locator}|$.

Annotation. Where it is not obvious which object is being represented in a value storage structure (for example, because it is represented via an indirection) the text of the corresponding object from the canonical form of the update scheme or archetype definition from which it is taken will be included, in italics, under its representation. A pointer to an object will be indicated by prefixing an indirection symbol ($*$) to the object’s text. See example 1(b) for an example.

Standard Environment. This thesis does not specify the standard environment of Update Plans. Where necessary, requirements for the standard environment will be indicated.

Memory Access. As mentioned in chapter five, certain mappings must be defined in the standard environment. One of these is the memory access function. The ‘@’ symbol will be used for the memory access function — i.e. $@(l, n)$ will, when applied to a configuration containing $l[c]l + n$, yield c . The default value of n is \mathbb{L} , so that $@(l, \mathbb{L})$ may be written $@l$. As usual, the @ operator has a higher priority than + or –, so that $@Cell + 3 = (@Cell) + 3$. Again @ can be defined as an archetype.

$$@(\mathbf{l}, \mathbf{n}) = \mathbf{val} \mid \mathbf{l}[\mathbf{val}]\mathbf{l} + \mathbf{n}.$$

Expression Evaluation. There must be some internal representation of expressions, and some mechanism for evaluating them. The presentation of the implementation must distinguish between the text of expressions, as they appear in update schemes, and the internal representation of expressions. As usual, a typewriter font will be used for the text as it appears in an update scheme. An italic typewriter font will be used to indicate an internal representation of an expression. The evaluation function will be written $\mathcal{E}[\cdot]$. E.g. the internal representation of the expression $@Cell + 3$ is $@Cell + 3$. Given a configuration satisfying $Cell[4]$ the value of $\mathcal{E}[@Cell + 3]$ will be $@Cell + 3$, or 7, the representation of which is 7. It is assumed that the length of an object can be determined by an examination of its internal representation.

Uninstantiated Values. It will be assumed that there are (possibly typed) special constants representing uninstantiated values. The “don’t care” symbol ‘_’ will be used for these constants.

2 Basic Update Schemes

This section defines a straightforward implementation of basic Update Plans. This will be extended in section 3 to cover archetypes. In both this

	Stores		Registers
Values	Storage for fixed length values and pointers to values	V	Pointer to Values
Heap	Storage for other values	H	Pointer to Heap
Rhs	Storage for the internal representations of right hand sides	R	Pointer to Rhs
Guard	Storage for internal representation of guards	G	Pointer to Guard
NO	Register holding the number of the current scheme		

Figure 1: Stores and registers required for the implementation of basic Update Plans.

section and section 3 the update schemes and archetypes will be considered to be deterministic. In particular, it will be assumed that applicability of archetypes can be decided at the time of their expansion. Nondeterminism and backtracking will be introduced in section 4, and the backtracking mechanism adapted to implement parallel blocks in section 4.3.

Implementing basic Update Plans is simple. Each update scheme is numbered, and a register, **NO**, is added to the machine, in which the number of the update scheme currently under consideration is stored. If the scheme can be applied it is, and the counter in **NO** is reset to 1. If the counter passes the number of schemes in the plan, then no applicable scheme could be found, and execution must stop. This is expressed in the following update schemes, which replace the n^{th} update scheme, $lhs_n = [guard_n] \Rightarrow rhs_n$. N is the number of update schemes in the update plan.

$$\begin{aligned} \mathbf{NO}[\mathbf{no}] \quad lhs_n &= [guard_n] \Rightarrow \mathbf{NO}[1] \quad rhs_n; \\ \mathbf{NO}[\mathbf{no}] &= [n \leq N] \Rightarrow \mathbf{NO}[\mathbf{no} + 1]. \end{aligned}$$

(Note the use of alternatives in this scheme.) This in no way changes the semantics of the update plan, it merely makes the sequential consideration of update schemes explicit. A fuller specification of an implementation is given in section 2.2.

2.1 Internal Representation

This section defines, mostly by archetype, the internal representation of update schemes. The stores and registers required by the representations are given in figure 1.

Values. In order to make values accessible they are classified as being either the value of a variable of static length — i.e. one with a length determinable from the update scheme and typing information alone —

or one of a variable of dynamic length. The former are stored in **Values** while the latter are placed in **Heap**. The values in **Heap** are accessed via pointers in **Values**. An archetype can now be defined for each update scheme, which creates a list of the values appearing in the update scheme.

Example 1 (b)

The following archetype defines the structure required for storage of the values in the **NEW_ENV** update scheme, where no is the number assigned to this update scheme.

$$\begin{aligned}
& \mathbf{vals}() = \\
& \quad \mathbf{NO}[no] \mathbf{V}[\mathbf{val}] \mathbf{H}[\mathbf{hp}] \\
\Rightarrow & \\
& \quad \mathbf{V}[\mathbf{val}'] \mathbf{H}[\mathbf{hp}'] \\
& \quad \mathbf{hp}[\underbrace{\mathbf{@}(\mathbf{@ENV}, \mathbf{@}(\mathbf{@PC} + |\mathbf{NEW_ENV}|, |\mathbf{Num}|))}_{env}] \mathbf{hp}' \\
& \quad \mathbf{val}[\underbrace{\mathbf{PC}}_{PC} \underbrace{\mathbf{@PC}}_{x_0} \underbrace{\mathbf{PC} + \mathbf{IL}}_{x_1} \underbrace{\mathbf{@}(\mathbf{@PC}, |\mathbf{NEW_ENV}|)}_{new_env} \\
& \quad \quad \underbrace{\mathbf{@PC} + |\mathbf{NEW_ENV}|}_{x_2} \underbrace{\mathbf{@}(\mathbf{@PC} + |\mathbf{NEW_ENV}|, |\mathbf{Num}|)}_n \\
& \quad \quad \underbrace{\mathbf{@PC} + |\mathbf{NEW_ENV}| + |\mathbf{Num}|}_{x_3} \underbrace{\mathbf{HEAP}}_{HEAP} \underbrace{\mathbf{@HEAP}}_h \\
& \quad \quad \underbrace{\mathbf{HEAP} + \mathbf{IL}}_{x_4} \underbrace{\mathbf{ENV}}_{ENV} \underbrace{\mathbf{@ENV}}_e \underbrace{\mathbf{ENV} + \mathbf{IL}}_{x_5} \underbrace{\mathbf{hp}}_{*env} \\
& \quad \quad \underbrace{\mathbf{@ENV} + \mathbf{@}(\mathbf{@PC} + |\mathbf{NEW_ENV}|, |\mathbf{Num}|)}_{e+n} \underbrace{\mathbf{PC} + \mathbf{IL}}_{x_6} \\
& \quad \quad \underbrace{\mathbf{ENV} + \mathbf{IL}}_{x_7} \underbrace{\mathbf{@HEAP} + \mathbf{@}(\mathbf{@PC} + |\mathbf{NEW_ENV}|, |\mathbf{Num}|)}_i \\
& \quad \quad \underbrace{\mathbf{HEAP} + \mathbf{IL}}_{x_8} \mathbf{val}'.
\end{aligned}$$

The expressions for the values of the terms in the update scheme can be derived during grounding analysis.

Note the indirection, via **hp**, to **env**.

A practical implementation would almost certainly instantiate these values one by one, in the order of their dependencies, allowing subexpressions such as **@PC** to be shared. In fact such values would probably not be stored in **Values**, but in a temporary register. The structure presented here should be taken as symbolising internal storage (and access) of values, and not as a definition of a suitable structure for a practical implementation.

Configurations. A locator expression is stored as a triple, containing pointers to its left locator, contents and right locator.

The right hand side of the `NEW_ENV` update scheme could be defined by the archetype:

$$\begin{aligned}
 \text{rhs}() &= \\
 &\text{NO}[no] \text{R}[\text{rhs}] \text{val}^{\text{env}}[\text{hp}] \\
 \implies & \\
 &\text{R}[\text{rhs}'] \\
 &\text{rhs}[\text{val}^{\text{PC}} \quad \text{val}^{\text{x3}} \quad \text{val}^{\text{x6}} \\
 &\quad \text{val}^{\text{ENV}} \quad \text{val}^{\text{h}} \quad \text{val}^{\text{x7}} \\
 &\quad \text{val}^{\text{h}} \quad \text{hp} \quad \text{val}^{\text{i}} \\
 &\quad \text{val}^{\text{HEAP}} \quad \text{val}^{\text{i}} \quad \text{val}^{\text{x8}}] \text{rhs}'.
 \end{aligned}$$

Guards. An internal representation for guards must also be available. Precise details will not be given here, since this would entail a full specification of function representation and evaluation in the standard environment. A guard will be assumed to consist of a conjunction of simpler conditions that can be evaluated by some mechanism defined in the standard environment. A guard is then translated to a list of pointers to internal representations of its constituent clauses.

The guard of the `NEW_ENV` update scheme contains only one clause. The archetype defining the internal representation of this guard is:

$$\begin{aligned}
 \text{grd}(\text{grd}, \text{grd}', \text{val}) &= \\
 &\text{NO}[no] \text{G}[\text{grd}] \text{H}[\text{hp}] \text{val}^{\text{new-env}}[\text{new_env}] \\
 \implies & \\
 &\text{G}[\text{grd}'] \text{H}[\text{hp}'] \text{grd}[\text{hp}]\text{grd}' \\
 &\text{hp}[\text{new_env} = \text{NEW_ENV}]\text{hp}'.
 \end{aligned}$$

2.2 An Implementation

The specification is given in terms of phases of an “instruction cycle”, in a manner similar to the update schemes in section 5.3 of chapter seven. Each phase may, itself, consist of many steps. The phases are *creation*, *probation*, *verification* and *application*. The archetypes used here are those defined above.

Creation. The creation phase creates the internal representation of the scheme. The constant G_0 is the base value of G — i.e. the value G contained in the initial configuration. The creation phase also increments the contents of `NO`.

[2.1] CREATE `NO[no]` $= [n \leq N] \Rightarrow$
 TRY G_0 `vals()` `grd()` `rhs()` `NO[no + 1]`.

Probation. Since the update scheme is in canonical form applicability can be checked simply by evaluating the guard. This is done by “walking through” the internal representation of the guard, evaluating its constituent clauses. If all components of the guard have been checked the scheme is applicable. If the current component is true, then the following component must be checked. If it is false then the guard is false and the next scheme must be tried.

$$\begin{array}{lcl}
[2.2] \text{ TRY } \text{grd}' \text{ G}[\text{grd}'] & \implies & \text{VERIFY } R_0 \text{ } R_0; \\
\text{TRY } \text{grd} \text{ grd}[\text{h}]\text{grd}' \text{ h}[\text{cond}] & =[\mathcal{E}[\text{cond}]] \Rightarrow & \text{TRY } \text{grd}'; \\
\text{"} & \implies & \text{CREATE.}
\end{array}$$

Verification. The right hand side must be checked for consistency. This is done by taking each locator expression in turn and checking it for consistency with every other locator expression in the right hand side. The first argument of VERIFY is the left locator of the locator expression being checked, the second is that of the locator expression it is being checked against.

The update schemes in this paragraph are a series of alternatives. Commentary has been added to make the intention clearer.

The end of the right hand side has been reached. It is consistent and may be applied.

$$[2.3] \text{ VERIFY } \text{rhs}' \text{ rhs}' \text{ R}[\text{rhs}'] \implies \text{APPLY } R_0;$$

The locator expression under consideration has been found to be consistent with the remainder of the right hand side. Move on to the next one.

$$\begin{array}{l}
[2.4] \text{ VERIFY } \text{rhs}_1 \text{ rhs}' \text{ R}[\text{rhs}'] \text{ rhs}_1[\text{l}_1 \text{ c}_1 \text{ r}_1]\text{rhs}'_1 \\
\implies \text{VERIFY } \text{rhs}'_1 \text{ rhs}'_1;
\end{array}$$

The two locator expressions do not overlap and are therefore consistent. Continue verification.

$$\begin{array}{l}
[2.5] \text{ VERIFY } \text{rhs}_1 \text{ rhs}_2 \text{ rhs}_1[\text{v}_{1_1} \text{ v}_{c_1} \text{ v}_{r_1}] \text{ rhs}_2[\text{v}_{1_2} \text{ v}_{c_2} \text{ v}_{r_2}]\text{rhs}'_2 \\
\text{v}_{1_1}[\text{l}_1] \text{ v}_{r_1}[\text{r}_1] \quad \text{v}_{1_2}[\text{l}_2] \text{ v}_{r_2}[\text{r}_2] \\
= [\text{r}_1 \leq \text{l}_2 \vee \text{r}_2 \leq \text{l}_1] \Rightarrow \text{VERIFY } \text{rhs}_1 \text{ rhs}'_2;
\end{array}$$

The overlap, but the contents of the overlapping cells are consistent.

$$\begin{array}{l}
[2.6] \text{ VERIFY } \text{rhs}_1 \text{ rhs}_2 \text{ rhs}_1[\text{v}_{1_1} \text{ v}_{c_1} \text{ v}_{r_1}] \text{ rhs}_2[\text{v}_{1_2} \text{ v}_{c_2} \text{ v}_{r_2}]\text{rhs}'_2 \\
\text{v}_{1_1}[\text{l}_1] \text{ v}_{r_1}[\text{r}_1] \quad \text{v}_{1_2}[\text{l}_2] \text{ v}_{r_2}[\text{r}_2] \\
\text{v}_{c_1} + (\max(\text{l}_1, \text{l}_2) - \text{l}_1)[\text{c}]\text{v}_{c_1} + (\min(\text{r}_1, \text{r}_2) - \text{l}_1) \\
\text{v}_{c_2} + (\max(\text{l}_1, \text{l}_2) - \text{l}_2)[\text{c}]\text{v}_{c_2} + (\min(\text{r}_1, \text{r}_2) - \text{l}_2) \\
\implies \text{VERIFY } \text{rhs}_1 \text{ rhs}'_2;
\end{array}$$

The locator expressions are inconsistent. Abort this scheme and try the next one.

$$[2.7] \text{ VERIFY } \text{rhs}_1 \text{ rhs}_2 \implies \text{CREATE.}$$

Application. APPLY takes each locator expression in turn from the internal representation and translates it to a concrete configuration. It terminates when it has run through the complete right hand side, and

	Stores		Registers
Archetypes	Storage for pointers to as yet unexpanded archetypes	A	Pointer to Archetypes
Parameters	Storage for pointers to archetype parameters	P	Pointer to Parameters
Lhs	Storage for the internal representations of left hand sides	L	Pointer to Lhs
Temp	Temporary storage, in which the left locators of values provisionally instantiated are noted	T	Pointer to Temp
F	A frame register, in which the previous contents of the other registers are temporarily stored		

Figure 2: Additional stores and registers for the implementation of archetypes

resets the update scheme number to 1.

[2.8] `APPLY rhs' R[rhs'] ==> CREATE NO[1];`

[2.9] `APPLY rhs rhs[v1 vc vr]rhs' v1[l] vc[c] vr[r]
==> APPLY rhs' l[c]r.`

3 Archetypes

Not only can command style update schemes be considered to be the bodies of an archetype definition, as shown in chapter seven, all canonical update schemes can. The “root” archetype then has empty expansions, the whole update scheme appearing in the context. In implementing archetypes all the update schemes in an update plan will be considered to define such an archetype, here called **root**. This archetype has only the implicit parameters $[l, l]$, $[r$ and $r]$, and these are irrelevant, since the expansions are empty. An application of update scheme n can now be treated as an application of the n^{th} definition of **root**.

Some additional stores and registers are required for the implementation of archetypes. These are listed in figure 2.

3.1 Structures

Some changes need to be made to the structures defined in section 2.1, and some new structures need to be defined. As the `NEW_ENV` update scheme was the running example in section 2 so will the archetype definition in

example 2(a) be for this section.

————— Example 2 (a) —————

The following archetype definition is designed to illustrate the implementation of archetypes, and is not necessarily a definition of a “useful” archetype.

$$\begin{aligned} \text{expr}(l + r) &= \\ &\quad \text{a PLUS b a}[\text{expr}_1(l)] \text{ b}[\text{expr}_2(r)] \\ \implies & \\ &\quad \text{PLUS expr}_1(l) \text{ expr}_2(r). \end{aligned}$$

The canonical form of this definition is

$$\begin{aligned} \text{expr}([l, l], [r, r], l + r) &= \\ [l[\text{a}]x_0 x_0[\text{plus}]x_1 x_1[\text{b}]l] & \\ \text{a}[\text{expr}_1(\text{a}, x_2, x_4, x_5, l)]x_2 \text{ b}[\text{expr}_2(\text{b}, x_3, x_5, r), r]x_3 & \\ =[\text{plus} = \text{PLUS}] \implies & \\ [r[\text{PLUS}]x_4 & \\ x_4[\text{expr}_1(\text{a}, x_2, x_4, x_5, l)]x_5 x_5[\text{expr}_2(\text{b}, x_3, x_5, r), r]r]. & \end{aligned}$$

Values. If there were no recursive archetypes full expansion of all archetypes would always terminate, and archetypes could therefore be expanded, in a macro-like manner, without reference to the current configuration, and no special mechanism would be required for their implementation. However, archetypes may be recursive, and consequently they must be expanded “on the fly”. This means that many values cannot be instantiated immediately, but only after archetypes have been expanded. Another consequence of the archetype mechanism is that values will be shared by way of archetype parameters. This sharing is implemented by adding an extra indirection. All values are now stored on **Heap**, and pointers to these values are maintained in **Values**. These aspects are illustrated in example 2(b), in which as yet uninstantiated values are indicated by the “don’t care” symbol ‘_’.

————— Example 2 (b) —————

A value storage archetype for archetype definition 2(a).

$$\begin{aligned} \text{vals}() &= \\ \text{NO}[\text{EXPR no}] \text{V}[\text{val}] \text{H}[\text{hp}] & \\ \implies & \\ \text{V}[\text{val}'] \text{H}[\text{hp}'] & \\ \text{val}[\underbrace{\quad}_l \underbrace{\quad}_{l_1} \underbrace{\quad}_{r_1} \underbrace{\quad}_{r_2} \underbrace{\quad}_{*l+r} \underbrace{\quad}_b] & \\ \underbrace{\quad}_{x_0} \underbrace{\quad}_{plus} \underbrace{\quad}_{x_1} \underbrace{\quad}_a \underbrace{\quad}_l \underbrace{\quad}_{x_2} & \\ \underbrace{\quad}_r \underbrace{\quad}_{x_3} \underbrace{\quad}_{*PLUS} \underbrace{\quad}_{x_4} \underbrace{\quad}_{x_5}] \text{val}' & \\ \text{hp}[\text{@@val}^l + \text{@@val}^r] \text{hp}_1[\text{PLUS}] \text{hp}'. & \end{aligned}$$

Note that the **NO** register now also includes an identification of the archetype.

Parameters. A new structure is introduced in which archetype parameters are stored. The first parameter list in the structure is that of the definition. This is followed by parameter lists of archetype calls in the body of the definition. The parameters are addressed through **Values**.

————— Example 2 (c) —————

Parameter storage for archetype definition 2(a). The first five entries in the structure are the parameters of the archetype definition, the next five those of the call **expr₁(l)**, and the remainder those of **expr₂(r)**.

```

pars() =
  NO[EXPR no] P[par]
⇒
  P[par']
  par[vall  vall  valr  valr  vall+r
     vala  valx2 valx4 valx5 vall
     valb  valx3 valx5 valr  valr]par'.

```

Left Hand Sides. It now becomes necessary to have an internal representation of the left hand side, since information contained within the left hand side is required to resolve terms. This representation contains two types of structures, one representing ordinary locator expressions, and the other archetype calls. The first field of each type of structure contains a constant identifying its type. The remaining three items of a “**LOC**” structure are pointers to the left locator, the contents, and the right locator respectively. The second item in an “**ARCH**” structure identifies the archetype in question. The third points to the call’s parameter list, and the fourth will point to the parameter list of the definition when the call is expanded.

————— Example 2 (d) —————

```

lhs() =
  NO[EXPR no] L[lhs]
⇒
  L[lhs']
  lhs[LOC  vall  vala  valx0
     LOC  valx0 valplus valx1
     LOC  valx1 valb  vall
     ARCH EXPR parexpr1  -
     ARCH EXPR parexpr2  _]lhs'.

```

Guards. Guards are implemented as in section 2.

```

grd() =
  NO[EXPR no] G[grd] H[hp] valplus[plus]
  ⇒
  G[grd'] H[hp'] hp[plus = PLUS]hp' grd[hp]grd'.

```

Right Hand Sides. The representation of right hand sides is similar to that of left hand sides. However there is no need to store archetype calls on the right hand side since these are fully documented in the representation of the left hand side. Consequently it is no longer necessary to indicate the type of the structure (locator expression or archetype call) being stored.

```

rhs() =
  NO[EXPR no] R[rhs]
  ⇒
  R[rhs']
  rhs[vall' valPLUS valx4]rhs'.

```

Archetype Calls. The order in which archetypes are expanded depends on the order in which their parameters contribute to grounding — it may be necessary to expand one archetype in order to ground terms necessary for the expansion of another. This order can be determined during grounding analysis.

```

calls() =
  NO[EXPR no] A[a]
  ⇒
  A[a']
  a'[lhsexpr1 lhsexpr2]a.

```

3.2 An Implementation

The following implementation of archetypes assumes that the update plans being implemented are fully deterministic — i.e. that at any point during archetype expansion sufficient information is available (a sufficient set of terms is instantiated) to determine uniquely, by inspection of archetype definitions, which archetype definition, if any, is applicable. Provisional expansion, until the probation phase, is used to make this inspection possible. Values that need to be restored to an uninstantiated value if expansion is unsuccessful are stored in **Temp**.

In keeping with the idea that all applications of update schemes are cases of archetype expansion the creation phase is renamed *expansion*.

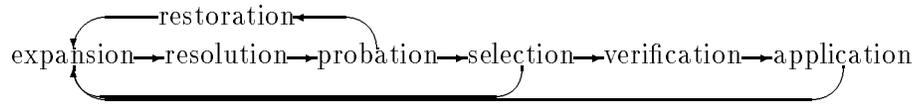


Figure 3: The instruction cycle for the archetype machine.

The initial configuration must satisfy $\text{NO}[\text{ROOT } 1]$, ensuring that expansion begins with the first update scheme in the update plan. Three new phases are needed in the “instruction cycle”, a *resolution phase* in which values are derived for as many terms from the current update scheme as possible, a *restoration phase* in which a provisional expansion is undone if probation is unsuccessful, and a *selection phase*, in which a choice is made between continuing expansion and application. The new instruction cycle is graphically represented in figure 3.

Expansion. The expansion phase constructs an expansion of the current archetype definition and provisionally appends this to the current scheme. The previous values of the pointers to the stores are temporarily stored in the frame register, allowing them to be restored if this expansion is unsuccessful.

[3.1] EXPAND L[lhs] G[grd] R[rhs] V[val] P[par] A[a] NO[arch no]
 $\Rightarrow [n \leq N^{\text{arch}}] \Rightarrow$
 RESOLVE vals() pars() calls() lhs() grd() rhs()
 F[lhs grd rhs val par a] NO[arch no + 1].

Resolution. The resolution phase is not presented in detail, involving as it does details of the expression evaluation mechanism from the standard environment. The resolution phase must:

- by inspection of archetype calls in the representation of the left hand side, match parameters of archetype calls and definitions,
- instantiate any expressions on the heap that become ground,
- examine the left and right hand side to determine if any other terms have become ground,
- note which terms have been instantiated during this resolution phase so that these instantiations can be undone if probation fails.

————— Example 2 (h) —————

Given the update scheme

$$\text{INF expr}(v) \implies \text{RP expr}(v).$$

or, in canonical form

$$\begin{aligned} & \text{PC}[\text{pc}]x_0 \text{pc}[\text{inf}]x_1 x_1[\text{expr}(x_1, x_2, x_5, x_2, v)]x_2 \\ & \Rightarrow [\text{inf} = \text{INF}] \Rightarrow \\ & \text{PC}[x_3]x_4 x_3[\text{RP}]x_5 x_5[\text{expr}(x_1, x_2, x_5, x_2, v)]x_2. \end{aligned}$$

the terms PC, pc, \mathbf{x}_0 , inf, \mathbf{x}_1 , \mathbf{x}_4 and RP will be instantiated before any archetype expansion takes place. Given the archetype definition from example 2(a)

$$\begin{aligned} \text{expr}(\{l, l\}, \{r, r\}, \mathbf{l} + \mathbf{r}) = \\ \{l[\mathbf{a}]\mathbf{x}_0 \mathbf{x}_0[\text{plus}]\mathbf{x}_1 \mathbf{x}_1[\mathbf{b}]\}l \\ \mathbf{a}[\text{expr}_1(\mathbf{a}, \mathbf{x}_2, \mathbf{x}_4, \mathbf{x}_5, \mathbf{l})]\mathbf{x}_2 \mathbf{b}[\text{expr}_2(\mathbf{b}, \mathbf{x}_3, \mathbf{x}_5, r, \mathbf{r})]\mathbf{x}_3 \\ =[\text{plus} = \text{PLUS}] \Rightarrow \\ \{r[\text{PLUS}]\mathbf{x}_4 \\ \mathbf{x}_4[\text{expr}_1(\mathbf{a}, \mathbf{x}_2, \mathbf{x}_4, \mathbf{x}_5, \mathbf{l})]\mathbf{x}_5 \mathbf{x}_5[\text{expr}_2(\mathbf{b}, \mathbf{x}_3, \mathbf{x}_5, r, \mathbf{r})]\}r\}. \end{aligned}$$

the first step of archetype expansion — i.e. before any further expansion of archetype calls in the archetype body takes place — will in turn instantiate $\{l, \mathbf{a}, \mathbf{x}_0, \mathbf{x}_1, \text{plus}, l\}$, \mathbf{b} and r from the archetype's body, and \mathbf{x}_2 from the update scheme in which it is called. Assuming that the left locator of the value storage structure of the update scheme is \mathbf{v}_0 , and that of the storage structure of the archetype \mathbf{v}_1 , immediately after the resolution phase of this cycle **Temp** will satisfy the configuration (with a slight abuse of notation)

$$\mathbf{T}[\mathbf{t}] \mathbf{t}[\mathbf{v}_0^{\mathbf{x}_2} \mathbf{v}_1^l \mathbf{v}_1^{\mathbf{a}} \mathbf{v}_1^{\mathbf{x}_0} \mathbf{v}_1^{\mathbf{x}_1} \mathbf{v}_1^{\text{plus}} \mathbf{v}_1^l \mathbf{v}_1^{\mathbf{b}} \mathbf{v}_1^r]\mathbf{T}_0$$

where \mathbf{T}_0 is the base value of **T** — i.e. the value in **T** in the initial configuration.

Probation. It is now possible, even likely, that terms in a guard are not instantiated. The evaluation mechanism must be able to handle uninstantiated values, returning a “don't care” value for any expression containing an uninstantiated value. Probation only fails if a condition is false. If the guard is non-false the current expansion is successful, and expansion may proceed. If probation fails the store pointers must be reset to their old values, and the restoration phase entered in order to undo any instantiations.

$$\begin{aligned} [3.2] \text{TRY } \text{grd}' \text{ G}[\text{grd}'] &\Rightarrow \text{SELECT}; \\ \text{TRY } \text{grd} \text{ grd}[\text{h}]\text{grd}' \text{ h}[\text{cond}] &= [\mathcal{E}[\text{cond} \neq \text{FALSE}]] \Rightarrow \text{TRY } \text{grd}'; \\ \text{TRY } \text{F}[\text{lhs} \text{ grd} \text{ rhs} \text{ val} \text{ par} \text{ a}] & \\ \Rightarrow \text{RESTORE } \text{L}[\text{lhs}] \text{ G}[\text{grd}] \text{ R}[\text{rhs}] \text{ V}[\text{val}] \text{ P}[\text{par}] \text{ A}[\text{a}]. & \end{aligned}$$

Restoration. The restoration phase must undo any instantiations from the preceding resolution phase.

$$\begin{aligned} [3.3] \text{RESTORE } \mathbf{T}[\mathbf{T}_0] &\Rightarrow \text{EXPAND}; \\ \text{RESTORE } \mathbf{T}[\mathbf{t}] \mathbf{t}[\mathbf{v}]\mathbf{t}' &\Rightarrow \text{RESTORE } \mathbf{T}[\mathbf{t}'] \mathbf{v}[_]. \end{aligned}$$

Selection. The selection phase inspects the archetype stack to determine if archetype expansion is complete. If A points to the bottom of the archetype stack, here indicated by the constant A_0 , then expansion is complete, and verification may commence. If not, the following archetype must be popped from the archetype stack, and the parameter list of the archetype's definition (which will be created by the expansion phase) connected to the archetype call on the left hand side.

```
[3.4] SELECT A[A0] ==> VERIFY R0 R0;
      SELECT A[a] a[l]a'[l][ARCH arch call _] P[def]
      ==> EXPAND A[a'] NO[arch l] l[ARCH arch call def].
```

Verification. Some indirections have been added to the internal representation of the right hand side. If these are taken into account the verification phase is identical to that in section 2. For example, update scheme 2.5 becomes

$$\begin{aligned} & \text{VERIFY rhs}_1 \text{ rhs}_2 \text{ rhs}_1[v_{1_1} \ v_{c_1} \ v_{r_1}] \text{ rhs}_2[v_{1_2} \ v_{c_2} \ v_{r_2}] \text{ rhs}'_2 \\ & \quad v_{1_1}[h_{1_1}] \ v_{r_1}[h_{r_1}] \ v_{1_2}[h_{1_2}] \ v_{r_2}[h_{r_2}] \\ & \quad h_{1_1}[l_1] \ h_{r_1}[r_1] \ h_{1_2}[l_2] \ h_{r_2}[r_2] \\ & \quad = [r_1 \leq l_2 \vee r_2 \leq l_1] \Rightarrow \text{VERIFY rhs}_1 \text{ rhs}'_2; \end{aligned}$$

Application. Application is again, when adapted to the extra indirections, more or less identical to the application phase of section 2. Upon completion, application must now call `EXPAND NO[ROOT 1]` rather than `CREATE NO[1]`.

```
APPLY rhs' R[rhs'] ==> EXPAND NO[ROOT 1];
```

4 Backtracking

Nondeterminism in update plans can have three sources. Firstly, more than one update scheme in the plan may be applicable. This is the simplest form to construct a backtracking mechanism for. Secondly, an update scheme containing semi-ground terms may have more than one applicable instantiation. Finally, more than one archetype expansion may be possible. An approach to backtracking across ambiguities of the first type is presented, in terms of the specification of basic Update Plans, in section 4.1. Adapting this to update plans with deterministic archetypes and no semi-ground terms would be simple. Section 4.2 discusses adapting the backtrack mechanism to cover semi-ground terms and nondeterministic archetypes. Finally section 4.3 shows how the backtrack mechanism can be used to implement parallel blocks.

4.1 Basic Backtracking

The symmetry of update schemes suggests that backtracking can be achieved by simply reversing update schemes — by swapping the left and right hand sides. Indeed this is almost the case. Reversing update schemes is the kernel of the backtracking mechanism. However, including the guard in the reversal is unnecessary, and in fact erroneous. Erroneous since a guard that was true before application is not necessarily true after application, and therefore before backtracking. Unnecessary because

backtracking is applied deterministically. The update scheme used for backtracking across update scheme $lhs = [guard] \Rightarrow rhs$ is $rhs \Rightarrow lhs$. This is called the *converse* of $lhs = [guard] \Rightarrow rhs$.

Example 3 (a)

The converse of the update scheme

$$[4.1] \text{ PC}[\text{pc}] \text{ pc}[\text{MOV } r_1 \ r_2] \text{ qc } r_1[\mathbf{x}] \Rightarrow \text{PC}[\text{qc}] \ r_2[\mathbf{x}].$$

is

$$[4.2] \text{ PC}[\text{qc}] \ r_2[\mathbf{x}] \Rightarrow \text{PC}[\text{pc}] \ \text{pc}[\text{MOV } r_1 \ r_2] \ \text{qc } r_1[\mathbf{x}].$$

The conceptual structure of backtracking is as follows. The numbering from section 2 is applied, using the register **NO**, and a backtrack stack, addressed by the register **B**, is introduced. Each update scheme is then replaced as follows. If the update scheme with index n is $lhs_n = [guard_n] \Rightarrow rhs_n$ then it is implemented in the backtracking version of the update plan as:

$$\begin{array}{l} \text{NO}[n] \ lhs_n^* \ \text{B}[\mathbf{c}] \ = [guard_n] \Rightarrow \ \text{NO}[1] \ rhs_n \ \text{B}[\mathbf{b}] \ \mathbf{b}[n \ \alpha_n] \ \mathbf{c}; \\ \text{NO}[n] \ \qquad \qquad \qquad \Rightarrow \ \text{NO}[\mathbf{n} + 1 \ \text{mod } \mathbf{N} + 1]. \end{array}$$

$$\text{NO}[0] \ rhs_n \ \text{B}[\mathbf{b}] \ \mathbf{b}[n \ \alpha_n] \ \mathbf{c} \ \Rightarrow \ \text{NO}[\mathbf{n} + 1 \ \text{mod } \mathbf{N} + 1] \ lhs_n^* \ \text{B}[\mathbf{c}].$$

where N is the number of update schemes in the update plan. The first scheme applies scheme n , if possible, and pushes any information necessary for backtracking onto the backtrack stack, and starts looking for schemes applicable to the new configuration by resetting **NO** to one. Its alternative simply moves on to the next possibility. The last scheme backtracks across an application of scheme n . The significance of the superscript $*$, and the definition of α_n , are given below. Informally the superscript indicates a simple transformation of the original update scheme, necessary to ensure correct backtracking, and α_n represents information that cannot be derived from the converse of the original update scheme, and that, therefore, must be preserved on the backtrack stack.

Completed Left Hand Sides. The superscript $*$ is used to indicate *completed left hand sides*. To ensure proper backtracking the left hand side of an update scheme must be extended to cover cells that are updated by the original update scheme but not specifically covered by the left hand side. This extension, in combination with pushing unprotected variables (see the next section), ensures that the contents of all cells will be properly restored. The left hand side is said to be *completed* with respect to the right hand side.

Example 3 (b)

The update — i.e. the unconditional update rule — specified by 4.1 is

$$[4.3] \Rightarrow \text{PC}[\text{qc}] \ r_2[\mathbf{x}].$$

(The variables appearing in 4.1 are here used to represent their values in the current configuration.) The update specified by 4.1's converse, 4.2, is

$$[4.4] \implies \text{PC}[\text{pc}] \text{pc}[\text{MOV } r_1 \ r_2] \text{qc } r_1[\mathbf{x}].$$

Combining these, i.e. applying first 4.1 and then 4.2, gives

$$\implies \text{PC}[\text{pc}] \text{pc}[\text{MOV } r_1 \ r_2] \text{qc } r_1[\mathbf{x}] \ r_2[\mathbf{x}],$$

which contains the locator expression $r_2[\mathbf{x}]$, which is not on the left hand side of 4.1 — i.e. the net effect of 4.3 and 4.4 on a configuration in which 4.1 is applicable is to update the contents of r_2 to \mathbf{x} .

Fortunately locator expressions such as $r_2[\mathbf{x}]$ in example 4(b) are easy to detect. Two locator expressions are said to *correspond* if they have the same left and right locators. Any locator expression on the right hand side of an update scheme without a corresponding locator expression on the left hand side will not be restored by backtracking using only converses. For each locator expression $l[x]r$ on the right hand side of an update scheme for which there is no corresponding expression on the left hand side, the locator expression $l[y]r$ must be added to the left hand side, where y is a new variable.

————— Example 3 (c) —————

Applying this extension to update scheme 4.1 gives

$$[4.5] \text{PC}[\text{pc}] \text{pc}[\text{MOV } r_1 \ r_2] \text{qc } r_1[\mathbf{x}] \ r_2[\mathbf{y}] \\ \implies \text{PC}[\text{qc}] \ r_2[\mathbf{x}].$$

A left hand side, *lhs*, completed in this way is written *lhs**. The fact that \mathbf{y} is non-ground in the converse of 4.5 will be addressed in the next section.

Unprotected Variables. Backtracking by application of a converse is only sure to work correctly if all terms appearing on the right hand side of the converse are fully ground. There is no guarantee that a semi-ground term will be restored to its original value. Indeed, as shown above, the converse may contain semi-ground terms or non-ground terms. Any semi-ground variable in a converse is said to be *unprotected*. Unprotected variables can easily be detected by the grounding mechanism.

————— Example 3 (d) —————

The converse of the extended scheme above is

$$\text{PC}[\text{qc}] \ r_2[\mathbf{x}] \\ \implies \text{PC}[\text{pc}] \ \text{pc}[\text{MOV } r_1 \ r_2] \ \text{qc } r_1[\mathbf{y}].$$

in which the variable \mathbf{y} on the right hand side is unground.

The simplest solution to this problem is to push any unprotected variables onto the backtrack stack. For each update scheme the sequence α_n is defined to be some fixed sequence containing exactly the unprotected variables of that update scheme.

The extended converse of an update scheme, combined with preservation of unprotected variables on the backtrack stack is sufficient to ensure correct backtracking. Backtracking can however be considerably simplified by application of the *backtrack rule*.

The Backtrack Rule. Completing left hand sides as described above is often unnecessary, causing the backtrack mechanism to preserve much superfluous information on the backtrack stack such as, for example, the value contained in the cell just above a stack when the stack is expanded by a push operation.

Example 4

The backtrack version of the update scheme

$$\text{PUSH } x \text{ SP}[t] \implies s[x]t \text{ SP}[s].$$

is

$$\begin{aligned} \text{PUSH } x \text{ NO}[no] \text{ SP}[t] \text{ s}[y]t \text{ B}[c] \\ \implies \text{NO}[1] \text{ s}[x]t \text{ SP}[s] \text{ B}[b] \text{ b}[n \ y]c. \end{aligned}$$

in which the value of y is, assuming that SP really does address a stack, unnecessarily pushed onto the backtrack stack.

A notational convention is needed to indicate changes that need not be reversed on backtracking. It is, of course, the responsibility of the update plan writer to ensure that the semantics of the plan will not be affected by use of this convention, though memory structure analysis as described in chapter five may be of some help. Meijer [53] introduced such a notational convention, known as the *backtrack rule*. The backtrack rule states

all cells that need to be restored upon backtracking are explicitly mentioned (covered) in the left hand side.

It is the task of the update plan writer to ensure that an update plan satisfies the backtrack rule. Application of the backtrack rule is equivalent to omitting the extension of the left hand side. In the specification of backtracking above any lhs^* may be replaced by lhs , to give a specification of backtracking with the backtrack rule applied.

4.2 Extended Backtracking

There is not a lot to be said about extending backtracking to cover semi-ground terms and ambiguous archetypes. In the first case since the ambiguity is at the level of the expression evaluation mechanism, and in the second since the implementation described here takes a brute force approach. In both cases the backtrack mechanism in section 4, adapted to the implementation in section 3 will be the departure point. In particular it will be assumed that a representation of both the left and right hand side of the update scheme being “undone” will have been constructed.

Semi-Ground Terms. The expression evaluation mechanism from the standard environment must be able to determine from examination of the values of semi-ground terms in the representation of the update scheme if all possible sets of values of have been tried. If all possible instantiations have been exhausted execution continues with the next update scheme. If not, the current update scheme (and archetype expansion) must be tried with the next possible instantiation.

If the values alone do not provide sufficient information for the backtracking mechanism of expression evaluation, sufficient data must be pushed to the backtrack stack to enable evaluation to recommence where it left off.

Ambiguous Archetypes. Unfortunately it is not clear how to implement efficient backtracking across ambiguous archetypes. A brute force approach is to reverse the complete expansion of the update scheme, and to determine the next applicable expansion, if any, from inspection of the expansion history which will at that point still be available immediately above the top of the backtrack stack, and to re-expand. An alternative approach is to treat each archetype definition as an independent update scheme, and to determine a minimal set of terms that guarantees grounding of the converse. This set of terms must then be pushed to the backtrack stack. It is a question for further investigation which of these approaches is the least inefficient.

4.3 Parallel Blocks

Parallel blocks can be implemented by “switching off” the application phase within a parallel block, and applying the techniques from the backtrack mechanism, within the parallel block, to construct a monolithic right hand side consisting of all the right hand sides of all applicable instantiations of update schemes within the parallel block. The same technique can be used on backtracking.

5 Conclusions

A possible implementation of update plans has been discussed. More efficient implementations are certainly possible. In particular, the mechanism for backtracking across archetypes is unsatisfactory in that a new expansion will repeat much of the work that has been undone in the backtracking stage. The efficiency of parallel blocks could also probably be greatly improved by sharing common substructures in the right hand side, again especially if the parallel block contains ambiguous archetypes. The greatest gain in efficiency, however, is in practice probably to be found in analysis of update schemes and archetypes to determine which terms contribute to the applicability of the schemes, and instantiating and checking these terms as soon as possible. In command style schemes, for example, it would be advisable to instantiate and check the value of the contents of the cell addressed by the programme counter before instantiating the rest of the scheme. The same technique can be applied to command archetypes, except that in this case it is the contents of the cell with left locator l that must be checked. This requires l to be

instantiated at the time, but in most applications of command archetypes
this will be the case.

In this thesis the syntax and semantics of Update Plans have been described. Some extensions to Update Plans (then called ‘Update Schemes’) as defined by Meijer [53] have been introduced. The thesis also presented some applications of Update Plans, both in specifying abstract and concrete machines, and in proving semantic properties of (programmes for) these machines. This chapter very briefly summarises the findings of this thesis, and presents some possible topics for future research.

1 Conclusions

Update Plans provide a readable and flexible specification formalism for low level languages. Their declarative style makes reasoning about such specifications relatively simple. The addition of the archetype mechanism greatly increases the expressive power of Update Plans, making a form of structured programming possible. The archetype mechanism also introduces new possibilities, such as simple specification of asynchronous parallel processors. Synchronous parallelism can also be specified by making use of the parallel block mechanism.

2 Further Research

Future research should take place on four fronts. Firstly, Update Plans should be placed in a wider theoretical context. Secondly, non-trivial applications should be undertaken, with an eye to developing pragmatics for update plan specifications. Thirdly, on a related note, Update Plans show promise as a didactic tool. This aspect should be further investigated. Finally a full implementation of Update Plans should be developed.

2.1 Theoretical Context

Rewrite Systems. Update Plans form an abstract rewrite system. Work on their relation to other ARS’s, and to graph rewrite systems in particular should lead to insights which would allow the well formedness conditions to be relaxed while still guaranteeing one or more of, for

example

- finite ambiguity
- the strong Church-Rosser property
- the weak Church-Rosser property
- modularity

Categories. Update Plans also define a category having consistent configurations as its objects, and update rules as its morphisms. Further investigation may make a category theoretical description of Update Plans possible.

2.2 Applications

Concrete Machines. The Update Plan specification formalism needs to be validated on a set of major sample applications, in order to develop the pragmatics of use. Among these applications are those in which *parallelism* plays a dominant rôle, e.g. the full specification of a “real” processor, such as SUN’s Sparc processor, Digital’s Alpha [2] or IBM’s PowerPC, or the (abstract) implementation of network protocols based on various communication primitives. Update Plans may also be applied to the specification of newer paradigms such as transport-triggered architectures [14, 30, 31, 32].

Abstract Machines. The archetype mechanism enables abstraction away from detail. An interesting challenge is to develop a suitable set of macros, possibly in combination with some simple transformations, such that the surface level language defined is a high level declarative formalism, for example the λ -calculus or some combinatory logic, while the fully expanded update schemes define a concrete implementation.

Hardware Specification. Update Plans could also be applied to the specification of hardware. Some preliminary work has been done in specifying elementary VLSI components [51].

Metrics. An annotation can be developed to indicate the cost of applying an update scheme or archetype. This would make it possible to apply Update Plans to the problem of compiler optimisation. The existence of a working implementation would make this even more useful, since the cost of proposed solutions could then be calculated by means of a simulation.

Formal Verification. For a specification and programming language to be useful at all it is imperative that one may prove a programme correct (e.g. equivalent to a specification) and/or derive a programme from a specification and/or prove that a specification or programme has certain desirable properties. Therefore a verification, and possibly transformation method should be developed, together with appropriate heuristic rules. Chapter six is only a first step in this direction.

2.3 Didactic Applications

The wide applicability and intuitive semantics of Update Plans make them well suited as a didactic tool for courses on machine architectures.

Update Plans should be described and explained in tutorials, containing many examples, in order to “market” and popularise them. A longer term project would be a survey of machine architectures presented in terms of update plans. Update Plans have been successfully used in the compiler construction course given at the University of Nijmegen [67], and in the machine architecture course at the University of Utrecht.

2.4 Implementation

A full implementation should be developed, embedded in an integrated collection of software, comprising at least a development environment, a compiler and a debugger.

The formalism should be further extended, in particular by introducing modules of some sort, possibly based on the archetype mechanism, in order to facilitate abstraction for information hiding, structure reuse, clarification, etc. Also, the parallel blocks introduced in chapter nine do not go much further than replacing nondeterministic execution of an applicable update scheme with simultaneous execution of all applicable update schemes. Such an approach will in practice, however, be too simple-minded. A host of artificial semaphore-like constructs would need to be introduced in order to maintain well-behaviour. Therefore it seems imperative to classify the different ways in which the concept of parallelism is used (e.g. co-operating processes, or systolic processes, or synchronous processes).

A major concern of computer science is the development of methods for proving software correct — either by verifying programmes already written or, preferably, by designing methodologies that ensure that only correct programmes are written. Much progress has been made, and since compilers themselves are programmes, application of the insights gained should make it possible to guarantee the correctness of compilers. Unfortunately compilers usually produce low level code, e.g. assembler, and there is seldom a formal definition of the semantics of such low level languages. Formal proofs cannot be based on informal descriptions. Update Plans are intended to fill this need. Update Plans constitute a formalism for the specification of low level languages. This is the *low level* part of the title of this thesis. The *high level* part refers to the formalism itself. Update Plans exhibit many of the features of other modern high level programming languages, giving them high abstractive power. Update Plans are a high level language for the specification of low level activities.

This thesis presents a complete formal definition of the syntax and semantics of Update Plans, and several examples of the formalism's use. It also presents a new typing system for Update Plans, and discusses the application of this to the detection of certain common types of memory use. These memory use paradigms have certain semantic properties, and this is also investigated. The thesis also introduces some extensions to Update Plans, in particular a macro-like mechanism, *archetypes*, and a degree of parallelism.

The archetype mechanism increases the expressive power of Update Plans, by significantly increasing their abstractive power, making it easier to write intuitive specifications. The archetype mechanism is illustrated by a specification of intermediate code such as may be emitted by a compiler before register allocation takes place. This is an interesting case in that the language specified is infinite, commands having a tree like structure in which subtrees represent the computation of intermediate results.

The addition of parallelism again extends the expressivity of Update

Plans to cover such features of low level architectures as pipelining. This is illustrated in a partial specification of the Berkeley RISC machine.

The formalism's usefulness, and popularity, will be increased by its realisation as a full programming language. Some aspects of such an implementation are also covered.

Further research into Update Plans could take place on the following fronts.

Update Plans must be developed into a full programming paradigm by the addition of some type of module mechanism.

A complete implementation of Update Plans, embedded in a software environment comprising supporting utilities, a (transformational) development tool, an interpreter and a debugger, should be developed.

The formalism should be further extended to cover, for example, different forms of parallelism.

The relation between Update Plans and (graph) rewrite systems should be investigated — in particular to what extent results from the latter “carry over” to the former.

Finally “real world” problems should be attacked. Full specifications of concrete machine architectures must be produced to illustrate the full power of the formalism. Related to this is work on Update Plans as a didactic tool. Course material should be written (informally) introducing Update Plans, and using them to introduce students to (various types of) machine architecture.

Een van de hoofddoelen van de informatica is het ontwikkelen van methoden om de correctheid van *software* te bewijzen — óf door verificatie van al geschreven programma's óf, nog beter, door het ontwerpen van methodieken die ertoe leiden dat alleen correcte programma's geschreven worden. Veel vooruitgang is al geboekt, en omdat *compilers* ook zelf programma's zijn kunnen deze methoden ook daarop toegepast worden. Dit zou het mogelijk moeten maken om ook de correctheid van *compilers* te bewijzen, maar helaas is het eindprodukt van een *compiler* meestal laag-nivo code, bijv. *assembler*, waarvan meestal geen formele specificatie voorhanden is. Formele bewijzen kunnen niet steunen op informele beschrijvingen. Update Plans zijn bedoeld om deze leemte te vullen. Update Plans zijn een formalisme voor het specificeren van laag-nivo programmeertalen. Dit is het *low level* deel van de titel van dit proefschrift. Het *high level* deel slaat op het formalisme zelf. Als programmeertaal hebben Update Plans veel gemeen met andere moderne hoog-nivo programmeertalen. Update Plans zijn een hoog-nivo taal voor het specificeren van laag-nivo activiteiten.

Dit proefschrift bevat een volledig formele definitie van de syntaxis en semantiek van Update Plans, geïllustreerd door enkele voorbeelden van hun toepassing. Een nieuw typeringssysteem wordt geïntroduceerd, en de toepassing daarvan m.b.t. het opsporen van bepaalde vaak voorkomende vormen van geheugengebruik wordt beschreven. Deze geheugengebruik paradigma's hebben bepaalde semantische eigenschappen, en dit wordt ook onderzocht. Ook beschrijft het proefschrift uitbreidingen van Update Plans, met name *archetypes*, een *macro*-achtig mechanisme, en een zekere mate van parallelisme.

Het *archetype* mechanisme verhoogt de uitdrukkingskracht van Update Plans door hun abstraherend vermogen significant te verhogen, waardoor het makkelijker wordt intuïtief specificaties te schrijven. Het mechanisme wordt geïllustreerd met een specificatie van tussencode zoals door een *compiler* geproduceerd zou kunnen worden, voordat *register allocation*

plaats vindt. Dit voorbeeld is interessant omdat de gespecificeerde taal oneindig is. Commando's hebben een boomachtige structuur waarin onderbomen het berekenen van tussenwaardes voorstellen.

Het toevoegen van parallellisme verhoogt wederom de uitdrukingskracht van Update Plans, die daardoor ook zulke eigenschappen van laag-nivo architecturen als *pipelining* kunnen beschrijven. Dit wordt geïllustreerd door een partiële specificatie van de Berkeley RISC machine.

De bruikbaarheid, en populariteit, van Update Plans zal verhoogd worden wanneer ze als volledige programmeertaal gerealiseerd worden. Enkele aspecten van zo'n implementatie worden ook behandeld.

Toekomstige onderzoek op het gebied van Update Plans zou als volgt plaats kunnen vinden.

Update Plans moeten verder ontwikkeld worden, tot een volledig programmeerparadigma, door het toevoegen van een module mechanisme.

Er moet een volledige implementatie van Update Plans komen, met ondersteunende faciliteiten, (transformationeel) ontwikkelingsgereedschap, een vertaler en een *debugger*.

Het formalisme moet verder uitgebreid worden, bijvoorbeeld om verschillende vormen van parallellisme te beschrijven.

Het verband tussen Update Plans en (graph-) herschrijfsystemen moet onderzocht worden — in het bijzonder de mate waarin resultaten in het tweede “geërfd” worden door het eerste.

Ten slotte moeten problemen uit de “echte wereld” aangepakt worden. Volledige specificaties van concrete machine architecturen zouden geschreven moeten worden om de volle kracht van Update Plans te illustreren. Hieraan verwant is aandacht voor Update Plans als didactisch middel. Collegedictaten zouden geschreven moeten worden waarin Update Plans (informeel) uitgelegd worden om vervolgens gebruikt te worden om studenten kennis te laten maken met machine-architecturen.

Curriculum Vitæ

19th January 1956 Date of birth. Ashbourne, UK.

4th July 1977 BSc Mathematics (Div. II), University of Manchester

28th April 1989 “Doctoraal” Computer Science (*cum laude*),
University of Nijmegen

1st May 1989 – 30th September 1991 Research Assistant, Esprit
Basic Research Action no. 3147, the *Phoenix* project, University
of Nijmegen

1st October 1991 – 16th December 1993 “Assistent in Opleiding”
(Research Assistant), University of Nijmegen

1st March 1994 – 31st July 1994 Systems Developer, University
of Nijmegen

1st October 1994 – 28th February 1995 Consultant, University
of Nijmegen

from 1st March 1995 Research Assistant, University of York

Other important dates:

- 18th October 1989
- 11th September 1991

The numbers between brackets are the reference numbers of the works cited, the numbers in parenthesis the numbers of the pages on which they are cited.

- [1]
(4) Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman.
Compilers, Principles, Techniques and Tools.
Addison-Wesley, 1986.
- [2]
(116) Communications of the ACM, 26(2),
February 1993.
(Special issue on the Alpha chip).
- [3]
(41) K.R. Apt and G.D. Plotkin.
A Cook's tour of countable nondeterminism.
In *Proceedings of the 8th. Colloquium on Automata, Languages and Programming*,
volume 115 of *Lecture Notes in Computer Science*. Springer Verlag, 1981.
S. Even and O. Kariv (editors).
- [4]
(41) R.J.R. Back.
A continuous semantics for unbounded
nondeterminism.
Theoretical Computer Science, 23:187–210,
1983.
- [5]
(1) R.C. Backhouse, P.J. de Bruin, G. Malcolm,
E. Voermans, and J. van der Woude.
Relational catamorphisms.
In B. Möller, editor, *Proceedings of the IFIP*

TC2 Working Conference on Constructing Programs from Specifications, pages 287–318, Amsterdam, 1991. North-Holland Publishing Company.

- [6]
(2) H.P. Barendregt, M.C.J.D. van Eekelen, J.R.W. Glauert, J.R. Kennaway, M.J. Plasmeijer, and M.R. Sleep.
Term graph rewriting.
In *Proceedings of Parallel Architectures and Languages in Europe (PARLE), Volume II: Parallel Languages*, volume 259 of *Lecture Notes in Computer Science*. Springer Verlag, Eindhoven, The Netherlands, 1987.
J.W. de Bakker, A.J. Nijman and P.C. Treleaven (editors).
- [7]
(1) F.L. Bauer, R. Berghammer, M. Broy, W. Dosch, F. Geiselbrechtinger, R. Gnatz, E. Hangel, W. Hesse, B. Krieg-Brückner, A. Laut, T. Matzner, B. Möller, F. Nickl, H. Partsch, P. Pepper, K. Samelson, M. Wirsing, and H. Wössner.
The Munich Project CIP. Volume I: The Wide Spectrum Language CIP-L, volume 183 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin/Heidelberg/New York, 1985.
- [8]
(1) F.L. Bauer and H. Wössner.
Algorithmic Language and Program Development.
Springer-Verlag, Berlin, 1982.
- [9]
(1) R.M. Burstall and J. Darlington.
A transformation system for developing recursive programs.
Journal of the ACM, 24(1):44–67, January 1977.
- [10]
(2) Volker Claus, Hartmut Ehrig, and Grzegorz Rozenburg, editors.
Graph Grammars and their Applications to Computer Science and Biology, volume 73 of *Lecture Notes in Computer Science*. Springer Verlag, 1979.
- [11]
(38) Implementation of CLEAN on MAC II-fx, see [72].

- [12]
(5) Todd A. Cook, Paul D. Franzon, Ed Harcourt, and Thomas K. Miller III.
System-level specification of instruction sets.
In *Proceedings of the 1993 IEEE International Conference on Computer Design*, pages 552–557, 1993.
- [13]
(5) Todd A. Cook and Ed Harcourt.
A functional specification language for instruction set architectures.
In *Proceedings of the IEEE International Conference on Computer Languages*, pages 11–19, 1994.
- [14]
(116) Henk Corporaal and Hans (J.M.) Mulder.
MOVE: A framework for high-performance processor design.
In *Supercomputing-91*, Albuquerque, New Mexico, USA, November 1991.
- [15]
(5) J.W. Davidson and C.W. Fraser.
The design and application of a retargetable peephole optimizer.
ACM Transactions on Programming Languages and Systems, 2(2):191–202, April 1980.
- [16]
(5) J.W. Davidson and C.W. Fraser.
Eliminating redundant object code.
In *Conference Record of the 9th Annual ACM Symposium on Principles of Programming Languages*, Albuquerque, New Mexico, USA, 1982.
- [17]
(30) N. G. de Bruijn.
Lambda calculus notations with nameless dummies, a tool for automatic formula manipulation.
Indagationes Mathematicæ, proceedings of the Koninklijke Nederlandse Akademie der Wetenschappen, pages 381–392, 1972.
- [18]
(55) Pierre Deransart, Martin Jourdan, and Bernard Lorho.
Attribute grammars : definitions, systems, and bibliography, volume 323 of *Lecture Notes in Computer Science*.
Springer-Verlag, Berlin, 1988.

- [19]
(4) Digital Equipment Corporation.
PDP-11 Processor Handbook, 1971.
- [20]
(1) Charles Donnelly and Richard Stallman.
*BISON: The YACC-compatible Parser
Generator*, 1988.
- [21]
(2) Hartmut Ehrig, Manfred Nagl, and
Grzegorz Rozenburg, editors.
*Graph-Grammars and their Applications to
Computer Science*, volume 153 of *Lecture
Notes in Computer Science*.
Springer Verlag, 1983.
- [22]
(4) Jon Fairbairn and Stuart Wray.
Tim: A simple, lazy abstract machine to
execute supercombinators.
In *Functional Programming Languages
and Computer Architecture*, volume 274
of *Lecture Notes in Computer Science*.
Springer Verlag, 1987.
Gilles Kahn (editor).
- [23]
(55) Gilberto Filè.
Theory of Attribute Grammars.
PhD thesis, Technische Hogeschool Twente,
1983.
- [24]
(69) H. Franzen and B. Hoffman.
Automatic determination of affix flow in
extended affix grammars.
In K. H. Böhling and P. P. Spies, editors,
Informatik-Fachberichte, Berlin, BRD,
1979. Gesellschaft für Informatik (GI),
Springer Verlag.
- [25]
(1) H. Franzen, B. Hoffman, and I.-R.
Petersen.
Ein Parser-Generator für erweiterte Affix-
Grammatiken.
Master's thesis, Technische Universität
Berlin, Fachbereich Informatik, October
1976.
Bericht Nr. 76-24.
- [26]
(5) C.W. Fraser.
A compact, machine-independent peephole
optimizer.
In *Conference Record of the 6th Annual*

ACM Symposium on Principles of Programming Languages, San Antonio, Texas, USA, 1979.

- [27]
(5) Robert Giegerich.
A formal framework for the derivation of machine-specific optimizers.
ACM Transactions on Programming Languages and Systems, 5(3):478–498, July 1983.
- [28]
(71) Robert Giegerich.
Implementierung von programmiersprachen. Technical report, Technische Fakultät, Universität Bielefeld, Postfach 86 40, 4800 Bielefeld 1, BRD, 1992.
Lecture notes for a course in compiler construction. (In German).
- [29]
(5) Ed Harcourt, Jon Mauney, and Todd Cook.
Functional specification and simulation of instruction set architectures.
In *Proceedings of the International Conference on Simulation and Hardware Description Languages*, pages 3–8. The Society for Computer Simulation, 1994.
- [30]
(116) Jan Hoogerbrugge.
Software pipelining for transport-triggered architectures.
Master’s thesis, Delft University of Technology, Delft, The Netherlands, December 1991.
- [31]
(116) Jan Hoogerbrugge and Henk Corporaal.
Comparing software pipelining for an operation-triggered and a transport-triggered architecture.
In *Compiler Construction: 4th International Conference, CC ‘92*, volume 641 of *Lecture Notes in Computer Science*. Springer Verlag, Paderborn, BRD, October 1992.
U. Kastens and P. Pfahler (editors).
- [32]
(116) Jan Hoogerbrugge, Henk Corporaal, and Hans Mulder.
Software pipelining for transport triggered

- architectures.
In *Proceedings of the 24th Annual Workshop on Microprogramming*, pages 74–81, Albuquerque, New Mexico, USA, November 1991.
- [33] S.C. Johnson.
(1) YACC - Yet Another Compiler-Compiler. Computer Science Technical Report No. 32, Bell Laboratories, Murray Hill, New Jersey, 1975.
- [34] U. Kastens, B. Hutt, and E. Zimmerman.
(1) *GAG: A Practical Compiler Generator*. Springer-Verlag, 1982.
- [35] Manolis G.H. Katevenis.
(90) *Reduced Instruction Set Architectures for VLSI*. ACM Doctoral Dissertation Awards. The MIT Press, 1985.
- [36] P. Klint.
(17,II.1) *The ASF+SDF Meta-Environment User's Guide, version 26*. Centrum voor Wiskunde en Informatica (CWI), February 1993.
Available by *ftp* from
`ftp.cwi.nl:pub/gipe/reports` as
`SysManual.ps.Z`.
- [37] P. Klint.
(17,II.1) A meta-environment for generating programming environments.
ACM Transactions on Software Engineering and Methodology, 2(2):176–201, 1993.
- [38] Donald E. Knuth.
(1) *The Art of Computer Programming*. Addison-Wesley, 1968.
- [39] Donald E. Knuth.
(55) Semantics of context-free languages.
Mathematical Systems Theory, 2(2):127–145, June 1968.
See also [40].
- [40] Donald E. Knuth.
(55,I.6) Semantics of context-free languages

- (correction).
Mathematical Systems Theory, 5(1):95–96,
May 1971.
- [41]
(4,38) Pieter Koopman.
*Functional Programs as Executable
Specifications.*
PhD thesis, University of Nijmegen, Toer-
nooiveld 1, Nijmegen, The Netherlands,
1990.
- [42]
(69) C. H. A. Koster.
Affix grammars.
In *Algol68 Implementation*. North Holland
Publishing Company, Utrecht, The
Netherlands, 1971.
- [43]
(1) C.H.A. Koster.
Using the CDL compiler compiler.
In F. L. Bauer and J. Eickel, editors,
*Compiler Construction, An Advanced
Course*, volume 21 of *Lecture Notes in
Computer Science*. Springer-Verlag, 1974.
- [44]
(55) P. Kühling.
*Affix-Grammatiken zur Beschreibung von
Programmiersprachen.*
PhD thesis, Technische Universität Berlin,
Fachbereich Informatik, February 1978.
In German.
- [45]
(40) Wil Lamain.
Update plans, implementatie aspecten.
Master's thesis, University of Nijmegen,
Toernooiveld 1, Nijmegen, The Nether-
lands, 1992.
(In Dutch).
- [46]
(56) B.M. Leavenworth.
Syntax macros and extended translation.
Communications of the ACM, 9(11):790–793,
November 1966.
- [47]
(vii,4) Hendrik C.R. Lock.
An abstract machine for the implementation
of functional logic programming languages.
Technical report, ESPRIT Basic Research
Action No. 3147 (the Phoenix Project),
1990.

- [48]
(vii) Hendrik C.R. Lock.
A specification of the JCode instructions
for the JUMP machine, 1991.
Personal communication.
- [49]
(55) Brian H. Mayoh.
Attribute grammars and mathematical
semantics.
SIAM Journal on Computing, 10(3):503–518,
August 1981.
- [50]
(1) E. Meijer, M.M. Fokkinga, and R. Pater-
son.
Functional programming with bananas,
lenses, envelopes and barbed wire.
In John Hughes, editor, *Functional Pro-
gramming and Computer Architecture*,
volume 523 of *Lecture Notes in Computer
Science*. Springer-Verlag, 1991.
- [51]
(1,116) Erik Meijer.
Calculating Compilers.
PhD thesis, University of Nijmegen, Toer-
nooiveld 1, Nijmegen, The Netherlands,
1992.
- [52]
(vii,25,29) Erik Meijer and Ross Patterson.
Down with lambda lifting!
Technical report, ESPRIT Basic Research
Action No. 3147 (the Phoenix Project),
1991.
- [53]
(vii,1,2,3,43,69,111,115) Hans Meijer.
Programmar: A Translator Generator.
PhD thesis, University of Nijmegen, Toer-
nooiveld 1, Nijmegen, The Netherlands,
1986.
- [54]
(1) P. Naur, (editor).
Revised report on the algorithmic language
algol 60.
Communications of the ACM, 6:1–17, 1963.
- [55]
(69) Hugh Osborne.
Choosy affix evaluation for EAG parsers.
Master's thesis, University of Nijmegen,
Toernooiveld 1, Nijmegen, The Nether-
lands, 1989.

- [56]
(vii) Hugh Osborne.
Update plans.
Periodica Polytechnica Ser. El. Eng.,
35(3):153–163, 1991.
- [57]
(vii) Hugh Osborne.
Update plans, an example: Flip.
Technical report, ESPRIT Basic Research
Action No. 3147 (the Phoenix Project),
1991.
- [58]
(vii) Hugh Osborne.
Update plans, an introduction.
Technical report, ESPRIT Basic Research
Action No. 3147 (the Phoenix Project),
1991.
- [59]
(vii) Hugh Osborne.
The semantics and syntax of update
schemes.
In *Code Generation — Concepts, Tools,
Techniques (Proceedings of the Interna-
tional Workshop on Code Generation,
Dagstuhl, Germany, 20–24 May 1991)*,
Workshops in Computing. Springer Verlag,
1992.
- [60]
(vii) Hugh Osborne.
Update plans.
In *Proceedings of the 25th Hawaii Interna-
tional Conference on System Sciences*.
IEEE Computer Society Press, 1992.
- [61]
(*stellingen*) Rohit J. Parikh.
On context-free languages.
*Journal of the Association for Computing
Machinery*, 13(4):570–581, October 1966.
- [62]
(1) H. Partsch.
*Specification and Transformation of
Programs - a Formal Approach to Software
Development*.
Springer-Verlag, Berlin, 1990.
- [63]
(4) Simon L. Peyton Jones and Jon Salkild.
The spineless tagless G-machine.
University College, London, 1989.
- [64]
(*stellingen*) D.L. Pilling.
Commutative regular equations and Parikh's

- theorem.
Journal of the London Mathematical Society,
2(6):663–666, 1973.
- [65] G.D. Plotkin.
(41) A powerdomain construction.
SIAM Journal on Computing, 5(3):452–487,
1976.
- [66] Sara and Tobias Osborne.
(vii) Personal communication.
- [67] Janos Sarbo.
(117) Collegedictaat vertalerbouw.
Technical report, University of Nijmegen,
Toernooiveld 1, Nijmegen, The Nether-
lands, 1991.
(In Dutch).
- [68] Daniel P. Siewiorek, C. Gordon Bell, and
(4) Allen Newell.
*Computer Structures: Principles and
Examples*.
Computer Science Series. McGraw-Hill,
1982.
- [69] C.P. Snow.
(1) *The Two Cultures and the Scientific
Revolution*.
Cambridge University Press, 1959.
- [70] STOP.
(1) *STOP International Summer School on
Constructive Algorithmics, Ameland*,
September 1989.
Lecture notes.
- [71] STOP.
(1) *Lecture Notes of the STOP 1992 Summer
School on Constructive Algorithmics,
Ameland*, 1992.
- [72] Marko van Eekelen.
(2,I.2) *Parallel Graph Rewriting*.
PhD thesis, University of Nijmegen, Toer-
nooiveld 1, Nijmegen, The Netherlands,
1988.
- [73] D. H. D. Warren.
(4) An abstract Prolog instruction set.

- [74] D. A. Watt.
(69) *Analysis Oriented Two Level Grammars*.
PhD thesis, University of Glasgow, 1974.
- [75] A. van Wijngaarden, B.J. Mailloux, J.E.L.
(1) Peck, C.H.A. Koster, M. Sintzoff, C.H.
Lindsey, L.G.L.T. Meertens, and R.G.
Fisker, editors.
*Revised Report on the Algorithmic Language
ALGOL 68*.
Springer-Verlag, 1976.

This appendix gives a full context free grammar for Update Plans as specified in this thesis. The grammar is given as an ASF+SDF [36, 37] specification, a specification formalism developed at the University of Amsterdam and the *Centrum voor Wiskunde en Informatica*. The specification can be read as a context free grammar, in which production rules read from right to left. The following notational conventions apply:

- S^* defines zero or more repetitions of sort S ;
- S^+ defines one or more repetitions of sort S ;
- $\{S \text{ sep}\}^*$ defines zero or more repetitions of sort S separated by the literal sep ;
- $\{S \text{ sep}\}^+$ defines one or more repetitions of sort S separated by the literal sep ;
- $[s]$ represents any one of the characters in the string s ;
- $\sim[s]$ represents any character not in the string s ;
- $\backslash \mathbf{t}$ is the horizontal tabulation character;
- $\backslash \mathbf{n}$ is the newline character.

The **{left}** and **{bracket}** annotation is applied to disambiguate parsings, as is the information provided under a **priorities** header.

module Plans

imports Layout Alternatives Archetypes ParallelBlocks Stores Types

exports

sorts SCRIPT PLAN ITEM

context-free functions

CONFIGURATION “.” PLAN → SCRIPT

ITEM* → PLAN

ALTERNATIVES → ITEM

ARCHETYPE-DEFINITION → ITEM

PARBLOCK → ITEM

STORE-DECLARATION → ITEM

TYPE-DECLARATION → ITEM

module ParallelBlocks

imports Alternatives

exports

sorts PARBLOCK

context-free functions

“(||” ALTERNATIVES* “|)” → PARBLOCK

module Alternatives

imports Schemes

exports

sorts ALTERNATIVES

context-free functions

{SCHEME “;”}+ “.” → ALTERNATIVES

module Archetypes

imports BasicArchetypes CommandArchetypes Terms

hiddens

sorts INDEX-OPT

exports

sorts ARCHETYPE-DEFINITION ARCHETYPE-CALL PARAMETERS

context-free functions

COMMAND-ARCHETYPE-DEFINITION → ARCHETYPE-DEFINITION

BASIC-ARCHETYPE-DEFINITION → ARCHETYPE-DEFINITION

ARCHETYPE-CALL → TERM

ARCHETYPE-NAME INDEX-OPT PARAMETERS → ARCHETYPE-CALL

“(” {TERM “,”}* “)” → PARAMETERS

INDEX → INDEX-OPT

→ INDEX-OPT

In the production rules for update schemes layout is forbidden between a locator and the '[' or ']' of the cell sequence of which it is a locator.

module Schemes

imports Terms

hiddens

sorts LOC-EXPR LEFT-SECTION LOCATOR

exports

sorts SCHEME CONFIGURATION TEXT CONTEXT REPEAT GUARD

context-free functions

CONFIGURATION GUARD CONFIGURATION → SCHEME

REPEAT GUARD CONFIGURATION → SCHEME

TEXT CONTEXT → CONFIGURATION

TERM * → TEXT

LOC-EXPR * → CONTEXT

“||” → REPEAT

“= [” TERM “] =>” → GUARD

“==>” → GUARD

LEFT-SECTION+ LOCATOR → LOC-EXPR

“?” TERM-OPT “[” TEXT “]” → LOC-EXPR

“!” TERM-OPT “[” TEXT “]” → LOC-EXPR

LOCATOR “[” TEXT “]” → LEFT-SECTION

TERM-OPT → LOCATOR

module BasicArchetypes

imports Terms Schemes

hiddens

sorts BASIC-DECLARATION BASIC-DEFINITION BASIC-BODY

exports

sorts BASIC-ARCHETYPE-DEFINITION

context-free functions

BASIC-DECLARATION BASIC-DEFINITION+ → BASIC-ARCHETYPE-DEFINITION

BASIC-ARCHETYPE-NAME PARAMETERS → BASIC-DECLARATION

"=" BASIC-BODY "." → BASIC-DEFINITION

CONFIGURATION GUARD CONFIGURATION → BASIC-BODY

REPEAT GUARD CONFIGURATION → BASIC-BODY

CONFIGURATION → BASIC-BODY

TEXT "|" CONTEXT GUARD CONTEXT → BASIC-BODY

TEXT "|" REPEAT GUARD CONTEXT → BASIC-BODY

TEXT "|" CONTEXT → BASIC-BODY

module CommandArchetypes

imports Terms Schemes

hiddens

sorts COMMAND-DECLARATION COMMAND-DEFINITION COMMAND-BODY

exports

sorts COMMAND-ARCHETYPE-DEFINITION

context-free functions

COMMAND-DECLARATION COMMAND-DEFINITION+ → COMMAND-ARCHETYPE-DEFINITION

COMMAND-ARCHETYPE-NAME PARAMETERS TEXT → COMMAND-DECLARATION

"=" COMMAND-BODY "." → COMMAND-DEFINITION

CONTEXT GUARD CONFIGURATION → COMMAND-BODY

REPEAT GUARD CONFIGURATION → COMMAND-BODY

CONTEXT → COMMAND-BODY

module Stores

imports Lexicon

exports

sorts STORE STORE-STRUCTURE STORE-DECLARATION

context-free functions

STORE-NAME	→ STORE-STRUCTURE
“(” STORE-STRUCTURE “)”	→ STORE-STRUCTURE { bracket }
STORE-STRUCTURE “++” STORE-STRUCTURE	→ STORE-STRUCTURE { left }
STORE-STRUCTURE “ ” STORE-STRUCTURE	→ STORE-STRUCTURE { left }
STORE-STRUCTURE “*”	→ STORE-STRUCTURE
“{” {STORE “,”}* “}”	→ STORE-DECLARATION
STORE-NAME	→ STORE
STORE-NAME “=” STORE-STRUCTURE	→ STORE

priorities

STORE-STRUCTURE “*”	→ STORE-STRUCTURE >
STORE-STRUCTURE “++” STORE-STRUCTURE	→ STORE-STRUCTURE >
STORE-STRUCTURE “ ” STORE-STRUCTURE	→ STORE-STRUCTURE

module Types

imports Terms Stores

exports

sorts TYPE-DECLARATION

context-free functions

{TERM “,”}* “:.” STORE-STRUCTURE “.”	→ TYPE-DECLARATION
--------------------------------------	--------------------

module Terms

imports BasicTerms Arithmetic Logic Ordering Memory Stores

exports

sorts TERM-OPT

context-free functions

“(” TERM “::” STORE-STRUCTURE “)” → TERM

TERM “*” → TERM

TERM “++” TERM → TERM {left}

TERM “|” TERM → TERM {left}

“(” TERM “)” → TERM {bracket}

TERM → TERM-OPT

→ TERM-OPT

priorities

TERM “*” → TERM > TERM “++” TERM → TERM > TERM “|” TERM → TERM

module BasicTerms

imports Lexicon

hiddens

sorts

exports

sorts TERM

context-free functions

VARIABLE → TERM

NUMBER → TERM

CHAR → TERM

SYMB-CONST → TERM

DONTCARE → TERM

module Arithmetic

imports BasicTerms

exports

context-free functions

TERM "+" TERM → TERM {left}

TERM "-" TERM → TERM {left}

TERM "×" TERM → TERM {left}

TERM "/" TERM → TERM {left}

TERM "^" TERM → TERM {left}

"-" TERM → TERM

priorities

"-" TERM → TERM > "^" > { "×", "/" } >

{ "+", TERM "-" TERM → TERM }

module Logic

imports BasicTerms

exports

context-free functions

TERM "^" TERM → TERM {left}

TERM "∨" TERM → TERM {left}

"¬" TERM → TERM

priorities

"¬" TERM → TERM > "^" > "∨"

module Ordering

imports BasicTerms

exports

context-free functions

TERM "<" TERM → TERM

TERM "≤" TERM → TERM

TERM "=" TERM → TERM

TERM "≠" TERM → TERM

TERM "≥" TERM → TERM

TERM ">" TERM → TERM

module Memory

imports Lexicon BasicTerms

exports

context-free functions

"|" STORE-NAME "|" → TERM

"@" "(" TERM "," TERM ")" → TERM

"@" TERM → TERM

module Lexicon

exports

sorts VARIABLE NUMBER CHAR SYMB-CONST DONTCARE
STORE-NAME
BASIC-ARCHETYPE-NAME COMMAND-ARCHETYPE-NAME ARCHETYPE-NAME
INDEX

lexical syntax

[0-9]+	→ NUMBER
“€” ~ [‘\n] “”	→ CHAR
[A-Z] [A-Z ‘0-9_]*	→ SYMB-CONST
[a-z] [a-z ‘0-9_]*	→ VARIABLE
“_”	→ DONTCARE
[A-Z] [A-Z0-9_’]* [a-z] [A-Za-z0-9_’]*	→ STORE-NAME
[a-z] [a-z ‘0-9_]*	→ BASIC-ARCHETYPE-NAME
[A-Z] [A-Z ‘0-9_]*	→ COMMAND-ARCHETYPE-NAME
BASIC-ARCHETYPE-NAME	→ ARCHETYPE-NAME
COMMAND-ARCHETYPE-NAME	→ ARCHETYPE-NAME
NUMBER	→ INDEX
“{” NUMBER “}”	→ INDEX

module Layout

exports

lexical syntax

[\t\n]*	→ LAYOUT
“ ”	→ LAYOUT

- [2.1] NEW_ENV n HEAP[h] ENV[e] $e[\text{env}]e+n$
 \implies ENV[h] $h[\text{env}]i$ HEAP[i].
- [2.2] BIND i r $r[h]$ ENV[e] $\implies e+i[h]$.
- [2.3] LOOKUP i r ENV[e] $e+i[h] \implies r[h]$.
- [2.4] PUSH r $r[h]$ SP[t] \implies SP[s] $s[h]t$.
- [2.5a] POP r SP[s] TP[u] $u[w] s[h]t$
 \implies $[s<w] \implies r[h]$ SP[t].
- [2.5b] POP PC[p] SP[s] TP[u] $u[s]v s[h]t$
 $h[\text{SUSP } c \ e] \text{ ENV}[f] \text{ HEAP}[i]$
 \implies PC[c] ENV[e] TP[v]
 $i[\text{SUSP } p \ f]j \text{ HEAP}[j] s[i]$.
- [2.6] SUSP r c ENV[e] HEAP[h]
 $\implies r[h] h[\text{SUSP } c \ e]i \text{ HEAP}[i]$.
- [2.7a] JUMP r $r[h]$ $h[\text{SUSP } c \ e] \implies$ PC[c] ENV[e].
- [2.7b] JUMP r $r[h]$ $h[\text{CONST}] \text{ TP}[u]$
 $u[s]v s[i]t i[\text{SUSP } c \ e]$
 \implies PC[c] ENV[e] TP[v] $s[h]$.
- [3.1] ASSIGN r $n \implies r[n]$.
- [3.2] CONST r $r[n]$ HEAP[h]
 $\implies h[\text{CONST } n]i r[h] \text{ HEAP}[i]$.

[4.1a] CALL r r[h] h[CONST] SP[t] \Rightarrow SP[s] s[h]t.

[4.1b] PC[p] p[CALL r]q r[h] h[SUSP c e]
ENV[f] HEAP[i] SP[t] TP[v]
 \Rightarrow PC[c] ENV[e] i[SUSP q f]j
HEAP[j] s[i]t SP[s] TP[u] u[s]v.

[5.1] ADD r1 r0 r1[h] h[CONST x] r0[i] i[CONST y]
 \Rightarrow r0[x+y].

[6.1] READ r ?[n] \Rightarrow r[n].

[6.2] WRITE r r[h] h[CONST n] \Rightarrow ![n].

Explanation of figure 1

1. If there is neither a read nor a write access the store can, for all intents and purposes, be said to not exist. However, no-access memories can appear as part of a configuration resulting from a transformation between configurations (see chapter eight).
2. Nothing is read from the store, and output is written sequentially, either right to left, or left to right.
3. This could be some sort of storage device, having data written to it. The data may be used in some other context, so that in the update plan under consideration it functions as “write-only memory”.
4. See output stream.
5. This case is further analysed in figure 2.
6. There seems to be no standard type of addressing structure corresponding to this combination.

read	write		
	no such access	structured access	unstructured access
no such access	no-access (1) memory	output (2) stream	write only (3) memory
structured access	input (4) stream	— (5)	??? (6)
unstructured access	read only (7) memory	heap (8)	general (9) memory

Figure 1: Access structure and store structure.

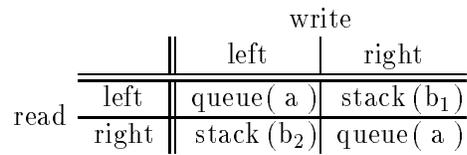


Figure 2: Structured read/write accesses.

7. For example, a programme store.
8. This case only has a real heap-like structure if an extra condition is guaranteed, namely that, at all times, any read address is in the luff of the write dedicated register. This is true if it can be shown that every locator in the store that is created, rather than copied, is, at the time of its creation, in the luff of the write dedicated register.
9. No structure can be detected.

Explanation of figure 2

- a) Two dedicated registers — one read register, one write register
- b) One register for read *and* write accesses, in case b₁ the stack grows to the right, in b₂ to the left.

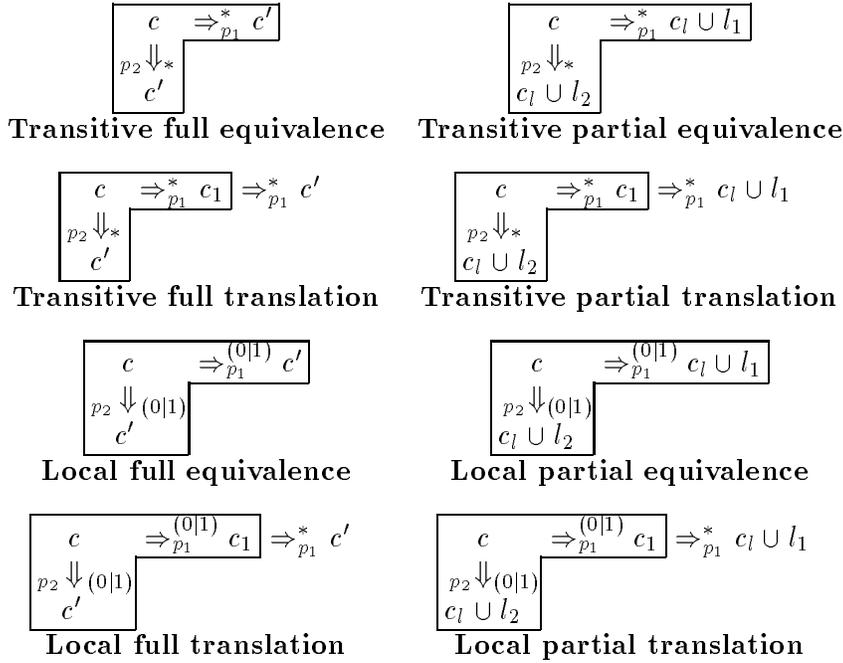


Figure 1: Equivalence and translation properties

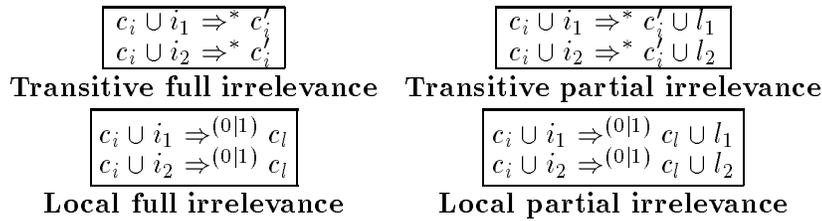


Figure 2: Irrelevance properties

$$\begin{aligned}
\text{MOV } \text{rop}(\mathbf{x}) \text{ wop}(\mathbf{a}, _) &\implies \mathbf{a}[\mathbf{x}]. \\
\text{MEA } \text{adr}(\mathbf{x}) \text{ wop}(\mathbf{a}, _) &\implies \mathbf{a}[\mathbf{x}]. \\
\\
\text{arith}(\mathbf{x}, \mathbf{y}, \mathbf{r}) \text{ rop}(\mathbf{x}) \text{ wop}(\mathbf{a}, \mathbf{y}) &\implies \mathbf{a}[\mathbf{r}]. \\
\\
\text{bool}(\mathbf{x}, \mathbf{y}, \mathbf{r}) \text{ rop}_1(\mathbf{x}) \text{ rop}_2(\mathbf{y}) &\implies \text{CC}[\mathbf{r}]. \\
\\
\text{PC}[\text{cp}] \text{ pc}[\text{jump}(\text{cond}) \text{ trg}(\mathbf{t})] \text{qc} &= [\text{cond}] \Rightarrow \text{PC}[\mathbf{t}]. \\
&= [\neg(\text{cond})] \Rightarrow \text{PC}[\text{qc}]. \\
\\
\text{PC}[\text{pc}] \text{ pc}[\text{JSR } \text{trg}(\mathbf{t})] \text{qc} \text{ SP}[\text{tp}] &\implies \text{PC}[\mathbf{t}] \text{ SP}[\text{sp}] \text{ sp}[\text{qc}] \text{tp}. \\
\\
\text{PC}[\text{pc}] \text{ pc}[\text{RET}] \text{ SP}[\text{sp}] \text{ sp}[\text{pc}] \text{tp} &\implies \text{PC}[\text{pc}] \text{ SP}[\text{tp}].
\end{aligned}$$

Figure 1: Tree and linear machine update schemes.

REG(r, b) r	= r[b].	
REGDEF(b, v) r	= r[b] b[v].	
BDISP(b+d, v) r d	= r[b] b+d[v].	
BDISPDEF(a, v) r d	= r[b] b+d[a] a[v].	
PREDEC(a, v) r	= r[b] a[v]b	\implies r[a].
POSTINC(b, v) r	= r[b] b[v]c	\implies r[c].

wop(a, v) = REG(a, v).	rop(v) = REG(., v).
= REGDEF(a, v).	= REGDEF(., v).
= BDISP(a, v).	= BDISP(., v).
= BDISPDEF(a, v).	= BDISPDEF(., v).
= PREDEC(a, v).	= POSTINC(., v).
	= IMM v.

adr(a) = REGDEF(a, .).	
= BDISP(a, .).	trg(t) = adr(t).
= BDISPDEF(a, .).	= LAB t.

arith(x, y, y+x) = ADD.	bool(x, y, x < y) = CLT.
arith(x, y, y - x) = SUB.	bool(x, y, x ≤ y) = CLE.
arith(x, y, y * x) = MUL.	bool(x, y, x = y) = CEQ.
arith(x, y, y/x) = DIV = [x ≠ 0]⇒ .	bool(x, y, x ≠ y) = CNE.
	bool(x, y, x ≥ y) = CGE.
	bool(x, y, x > y) = CGT.

jump(cc) = JT CC[cc].	
jump(¬(cc)) = JF CC[cc].	
jump(TRUE) = JMP.	

Figure 2: Linear machine archetypes.

$\text{REG}(r, b) \ r = r[b].$
 $\text{REGDEF}(b, v) \ r = r[b] \ b[v].$
 $\text{BDISP}(b+d, v) \ r \ d = r[b] \ b+d[v].$
 $\text{BDISPDEF}(a, v) \ r \ d = r[b] \ b+d[a] \ a[v].$
 $\text{PREDEC}(a, v) \ r = r[b] \ a[v]b \implies r[a].$
 $\text{POSTINC}(b, v) \ r = r[b] \ b[v]c \implies r[c].$
 $\text{IREG}(_, b) \ \text{ir}(b) = .$
 $\text{IREGDEF}(b, v) \ \text{ir}(b) = b[v].$
 $\text{IBDISP}(b+d, v) \ \text{ir}(b) \ d = b+d[v].$
 $\text{IBDISPDEF}(a, v) \ \text{ir}(b) \ d = b+d[a] \ a[v].$

$\text{kern}(a, v) = \text{REGDEF}(a, v). \quad \text{rop}(v) = \text{kern}(_, v).$
 $= \text{BDISP}(a, v). \quad = \text{IMM } v.$
 $= \text{BDISPDEF}(a, v). \quad = \text{REG}(_, v).$
 $= \text{IBDISP}(a, v). \quad = \text{POSTINC}(_, v).$
 $= \text{IBDISPDEF}(a, v). \quad = \text{IREG}(_, v).$
 $= \text{IREGDEF}(_, v).$

$\text{wop}(a, v) = \text{kern}(a, v). \quad \text{adr}(a) = \text{kern}(a, _).$
 $= \text{REG}(a, v). \quad \text{trg}(t) = \text{LAB } t.$
 $= \text{PREDEC}(a, v). \quad \text{tmp}(t) = \text{IREG}(_, t).$
 $= \text{IREGDEF}(a, v).$

$\text{iarith}(y+x) = \text{IADD } \text{rop}(x) \ \text{tmp}(y).$
 $\text{iarith}(y-x) = \text{ISUB } \text{rop}(x) \ \text{tmp}(y).$
 $\text{iarith}(y * x) = \text{IMUL } \text{rop}(x) \ \text{tmp}(y).$
 $\text{iarith}(y/x) = \text{IDIV } \text{rop}(x) \ \text{tmp}(y) = [x \neq 0] \Rightarrow .$

$\text{imov}(x) = \text{IMOV } \text{rop}(x) \ (\text{IREG NIL}). \quad \text{ir}(r) = \text{iarith}(r).$
 $\text{imov}(a) = \text{IMEA } \text{adr}(a) \ (\text{IREG NIL}). \quad = \text{imov}(r).$

$\text{arith}(x, y, y+x) = \text{ADD}. \quad \text{bool}(x, y, x < y) = \text{CLT}.$
 $\text{arith}(x, y, y-x) = \text{SUB}. \quad \text{bool}(x, y, x \leq y) = \text{CLE}.$
 $\text{arith}(x, y, y * x) = \text{MUL}. \quad \text{bool}(x, y, x = y) = \text{CEQ}.$
 $\text{arith}(x, y, y/x) = \text{DIV} = [x \neq 0] \Rightarrow . \quad \text{bool}(x, y, x \neq y) = \text{CNE}.$
 $\quad \text{bool}(x, y, x \geq y) = \text{CGE}.$
 $\text{jump}(cc) = \text{JT } \text{CC}[cc]. \quad \text{bool}(x, y, x > y) = \text{CGT}.$
 $\text{jump}(\neg(cc)) = \text{JF } \text{CC}[cc].$
 $\text{jump}(\text{TRUE}) = \text{JMP}.$

Figure 3: Tree machine archetypes.

$\text{MOV rop}(x) \text{ wop}(a, -) \implies a[x].$

$$\left[\begin{array}{l} \text{rop}(v_1) = \text{kern}(-1, v_1) \\ \text{wop}(a_2, v_2) = \text{IREGDEF}(a_2, v_2) \end{array} \right] \quad \begin{array}{l} x = v_1, a = a_2, \\ - = v_2 \end{array}$$

$\text{MOV kern}(-1, v_1) \text{ IREGDEF}(a_2, v_2) \implies a[x].$

$$\left[\begin{array}{l} \text{kern}(a_3, v_3) = \text{IBDISP}(a_3, v_3) \\ \text{IREGDEF}(a_4, v_4) \text{ ir}(a_4) = a_4[v_4] \end{array} \right] \quad \begin{array}{l} -1 = a_3, v_1 = v_3, \\ a_2 = a_4, v_2 = v_4 \end{array}$$

$\text{MOV IBDISP}(a_3, v_3) (\text{IREGDEF ir}(a_4))$

$$a_4[v_4] \implies a[x].$$

$$\left[\begin{array}{l} \text{IBDISP}(b_5+d_5, v_5) \text{ ir}(b_5) d_5 = b_5+d_5[v_5] \\ \text{ir}(r_6) = \text{imov}(r_6) \end{array} \right] \quad \begin{array}{l} a_3 = b_5 + d_5, \\ v_3 = v_5, a_4 = r_6 \end{array}$$

$\text{MOV (IBDISP ir}(b_5) d_5) (\text{IREGDEF imov}(r_6))$

$$b_5+d_5[v_5] \quad a_4[v_4] \implies a[x].$$

$$\left[\begin{array}{l} \text{ir}(r_7) = \text{imov}(r_7) \\ \text{imov}(x_8) = \text{IMOV rop}(x_8) (\text{IREG NIL}) \end{array} \right] \quad b_5 = r_7, r_6 = x_8$$

$\text{MOV (IBDISP imov}(r_7) d_5) (\text{IREGDEF (IMOV rop}(x_8) (\text{IREG NIL})))$

$$b_5+d_5[v_5] \quad a_4[v_4] \implies a[x].$$

$$\left[\begin{array}{l} \text{imov}(x_9) = \text{IMOV rop}(x_9) (\text{IREG NIL}) \\ \text{rop}(v_{10}) = \text{IMM } v_{10} \end{array} \right] \quad \begin{array}{l} r_7 = x_9, \\ x_8 = v_{10} \end{array}$$

$\text{MOV (IBDISP (IMOV rop}(x_9) (\text{IREG NIL})) d_5)$

$$(\text{IREGDEF (IMOV (IMM } v_{10}) (\text{IREG NIL})))$$

$$b_5+d_5[v_5] \quad a_4[v_4] \implies a[x].$$

$$\left[\text{rop}(v_{11}) = \text{kern}(-11, v_{11}) \right] \quad x_9 = v_{11}$$

$\text{MOV (IBDISP (IMOV kern}(-11, v_{11}) (\text{IREG NIL})) d_5)$

$$(\text{IREGDEF (IMOV (IMM } v_{10}) (\text{IREG NIL})))$$

$$b_5+d_5[v_5] \quad a_4[v_4] \implies a[x].$$

$$\left[\text{kern}(a_{12}, v_{12}) = \text{BDISP}(a_{12}, v_{12}) \right] \quad \begin{array}{l} -11 = a_{12}, \\ v_{11} = v_{12} \end{array}$$

$\text{MOV (IBDISP (IMOV BDISP}(a_{12}, v_{12}) (\text{IREG NIL})) d_5)$

$$(\text{IREGDEF (IMOV (IMM } v_{10}) (\text{IREG NIL})))$$

$$b_5+d_5[v_5] \quad a_4[v_4] \implies a[x].$$

$$\left[\text{BDISP}(b_{13}+d_{13}, v_{13}) r_{13} d_{13} = r_{13}[b_{13}] b_{13}+d_{13}[v_{13}] \right] \quad \begin{array}{l} a_{12} = b_{13} + d_{13}, \\ v_{12} = v_{13} \end{array}$$

$\text{MOV (IBDISP (IMOV (BDISP } r_{13} d_{13}) (\text{IREG NIL})) d_5)$

$$(\text{IREGDEF (IMOV (IMM } v_{10}) (\text{IREG NIL})))$$

$$r_{13}[b_{13}] b_{13}+d_{13}[v_{13}] b_5+d_5[v_5] \quad a_4[v_4] \implies a[x].$$

Figure 7: Textual expansion of archetypes in the tree machine

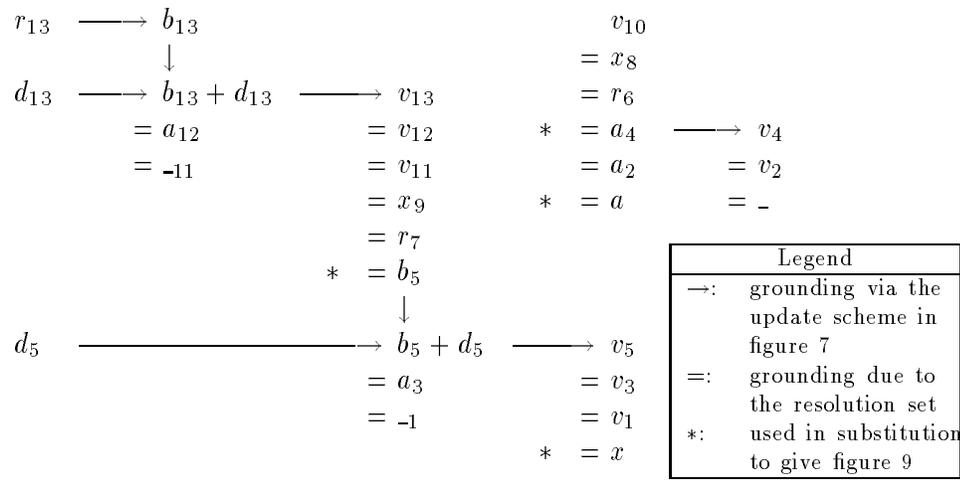


Figure 8: Parameter resolution for the archetype expansion in figure 7

```
MOV (IBDISP (IMOV (BDISP r13 d13) (IREG NIL)) d5)
  (IREGDEF (IMOV (IMM v10) (IREG NIL)))
r13[b13] b13+d13[v13] v13+d5[v5]      v10[v4]==>v10[v5].
```

Figure 9: The expanded update scheme from figures 7 and 8

$$\begin{aligned}
& \mathcal{T} \left[\left[\text{MOV (IBDISP (IMOV (BDISP r_{13} d_{13}) (IREG NIL)) d_5) (IREGDEF (IMOV (IMM v_{10}) (IREG NIL)))} \right] \right] n \\
= & \mathcal{IR}[\text{IREGDEF (IMOV (IMM v}_{10}) (IREG NIL))] n \\
& \mathcal{IR}[\text{IBDISP (IMOV (BDISP r}_{13} d_{13}) (IREG NIL)) d_5] n + 1 \\
& \text{MOV} (\mathcal{M}[\text{IBDISP (IMOV (BDISP r}_{13} d_{13}) (IREG NIL)) d_5] n + 1) \\
& \quad (\mathcal{M}[\text{IREGDEF (IMOV (IMM v}_{10}) (IREG NIL))] n) \\
= & \mathcal{T}[\text{IMOV (IMM v}_{10}) (IREG NIL)] n \\
& \mathcal{T}[\text{IMOV (BDISP r}_{13} d_{13}) (IREG NIL)] n + 1 \\
& \text{MOV (BDISP } n + 1 \text{ d}_5) (\text{REGDEF } n) \\
= & \mathcal{IR}[\text{IREG NIL}] n \\
& \mathcal{IR}[\text{IMM v}_{10}] n + 1 \\
& \text{MOV} (\mathcal{M}[\text{IMM v}_{10}] n + 1) (\mathcal{M}[\text{IREG NIL}] n) \\
& \mathcal{IR}[\text{IREG NIL}] n + 1 \\
& \mathcal{IR}[\text{BDISP r}_{13} d_{13}] n + 2 \\
& \text{MOV} (\mathcal{M}[\text{BDISP r}_{13} d_{13}] n + 2) (\mathcal{M}[\text{IREG NIL}] n + 1) \\
& \text{MOV (BDISP } n + 1 \text{ d}_5) (\text{REGDEF } n) \\
= & \mathcal{T}[\text{NIL}] n \\
& \text{MOV (IMM v}_{10}) (\text{REG } n) \\
& \mathcal{T}[\text{NIL}] n + 1 \\
& \text{MOV (BDISP r}_{13} d_{13}) (\text{REG } n + 1) \\
& \text{MOV (BDISP } n + 1 \text{ d}_5) (\text{REGDEF } n) \\
= & \text{MOV (IMM v}_{10}) (\text{REG } n) \\
& \text{MOV (BDISP r}_{13} d_{13}) (\text{REG } n + 1) \\
& \text{MOV (BDISP } n + 1 \text{ d}_5) (\text{REGDEF } n)
\end{aligned}$$

Figure 10: Transforming the tree code from figure 9 to linear code

$\text{reg}(\text{GBASE}+r) = \text{global}(r).$
 $\text{reg}(\text{cwp}+r) = \text{local}(r) \text{ CWP}[\text{cwp}].$
 $\text{global}(r) = r \text{ } \models [0 \leq r \leq 9] \Rightarrow .$
 $\text{local}(r) = r \text{ } \models [10 \leq r \leq 31] \Rightarrow .$
 $\text{rs}(0) = \text{reg}(\text{GBASE}).$
 $\text{rs}(\text{res}) = \text{reg}(\text{dst}) \text{ DST}[\text{dst}] \text{ RES}[\text{res}] \models [\text{dst} \neq \text{GBASE}] \Rightarrow .$
 $\text{rs}(\text{val}) = \text{reg}(\text{ea}) \text{ DST}[\text{dst}] \text{ ea}[\text{val}] \models [\text{ea} \neq \text{GBASE} \wedge \text{ea} \neq \text{dst}] \Rightarrow .$
 $\text{short}(\text{val}) = \text{REG rs}(\text{val}).$
 $\quad = \text{IMM val}.$
 $\text{scc}(_) = \text{OFF}.$
 $\text{scc}(\text{flags}) = \text{ON} \Rightarrow \text{CC}[\text{flags}].$
 $\text{aflg}(v) = (v <_s) (v = 0) \neg(\text{MIN}_s \leq_s v \leq_s \text{MAX}_s) (v >_u \text{MAX}_u) |.$
 $\text{arith}(x, y, x+y) = \text{ADD}.$
 $\text{arith}(x, y, x+y+c) = \text{ADDC } C[c].$
 $\text{arith}(x, y, x-y) = \text{SUB}.$
 $\text{arith}(x, y, x-y-(1-c)) = \text{SUBC } C[c].$
 $\text{arith}(x, y, y-x) = \text{SUBI}.$
 $\text{arith}(x, y, y-x-(1-c)) = \text{SUBCI } C[c].$
 $\text{arith}(x, y, r) \text{ scc}(\text{aflg}(r)) \text{ reg}(\text{ea}) \text{ rs}(x) \text{ short}(y)$
 $\quad \models [\text{ea} = \text{GBASE}] \Rightarrow .$
 $\quad " \models [\text{ea} \neq \text{GBASE}] \Rightarrow \text{ea}[r].$
 $\text{PC}[\text{pc}] \text{ pc}[\text{instruction}] \Rightarrow \text{IR}[\text{instruction}].$
 $\text{IR}[\text{arith}(x, y, r) \text{ scc}(\text{aflg}(r)) \text{ reg}(\text{dst}) \text{ rs}(x) \text{ short}(y)] \text{ PC}[\text{pc}]$
 $\quad \Rightarrow \text{RES}[r] \text{ DST}[\text{dst}] \text{ PC}[\text{pc}+\text{WORD}].$
 $\text{DST}[\text{dst}] \text{ RES}[v] \models [\text{dst} = \text{GBASE}] \Rightarrow .$
 $\quad " \models [\text{dst} \neq \text{GBASE}] \Rightarrow \text{dst}[v].$

The number(s) between brackets indicate(s) the page(s) on which the concepts are defined. A reference to another concept in the glossary is indicated by a bold font.

access (44)

A synonym for **locator expression**.

alternatives (16)

A series of **update schemes**. The first **applicable** scheme in the series is applied.

applicable (4,21)

An **update scheme** is applicable if its left hand side is consistent with the current **configuration**, and its **guard** evaluates to true.

application phase (101,108)

In an implementation, the phase in which an update scheme is applied.

archetype (55)

A macro like mechanism.

archetype expansion (60)

The mechanism by which **archetypes** are expanded to give **update schemes**. Also the **text** parts of an **archetype** body.

archetype grammar (65)

A context free grammar representing the structure of the **archetype** declarations.

archetype, left handed (59)

An **archetype** having an empty right hand side **expansion**.

archetype, right handed (59)

An **archetype** having an empty left hand side **expansion**.

archetype, ambidextrous (59)

A pair of **archetypes** of the same name, one **left handed**, the other **right handed**.

backtrack rule (111)

A convention for indicating which **cells** need not be restored on backtracking.

box diagrams (49)

A diagrammatical notation for existential dependencies.

canonical form (96)

A desugared **update scheme** in which all **guards** have been made explicit.

casting (40)

Forcing the value of a term to be of a different type than the type globally declared for that term.

cell (3,19)

An element of a **configuration** or **memory**.

command driven (14)

An **update plan** in which all **update schemes** are **commands**

command form (14)

A **configuration** containing a non-empty **command sequence**, or one in which the contents of the **register PC** are not specified.

command sequence (14)

A sequence of terms.

command (6,12,14)

An **update scheme** in which both the left and right hand side are in **command form**.

completed left hand sides (109)

An **update scheme** in which the left hand side is extended to cover all **cells** covered by the right hand side.

configuration (19)

A partial function from **locators** to values.

consistent (3,20)

A **configuration** is consistent if it does not specify conflicting contents for one and the same **cell**.

context (60)

The non-**text** part of a configuration, in particular in an **archetype** body.

converse (109)

The **update scheme** obtained by swapping the left and right hand sides of a given **update scheme**.

correspond (110)

Two **locator expressions** correspond if they have the same left and right **locators**.

creation phase (100)

In an implementation, the phase in which internal representations of an **update scheme** are created.

dedicated register (44)

A **register** is dedicated to a **store** if it only appears as a **locator** of an **access** of that store.

derivation (4)

The result of **applying** an **update plan** to a **configuration**.

development (23)

The result of repeated **applications** of an **update plan** to a **configuration**.

direct access (44)

An **access** with a **locator** in a **register**.

equivalence (47)

See **semantic equivalence**.

expansion (60)

See **archetype expansion**.

expansion phase (105,106)

In an implementation, the phase in which **archetype** calls are (textually) **expanded**.

final configuration (4)

A **configuration** to which none of the **update schemes** in an **update plan** are **applicable**.

final development (4,23)

A **development** which is a **final configuration**.

finite semi-constant (9)

A term representing a finite set of constants.

fully structured store (44)

A **store** that is both **read structured** and **write structured**, and in which **read accesses** and **write accesses** share the same **register**.

ground term (42)

A term for which a variable free value can be derived.

guard (3)

A condition for **applicability**.

initial configuration (4)

Specifies the initial state of the memory before any **update schemes** are applied.

lee (45)

Cells that have already been passed by the **dedicated register** of a **structured store**.

locator (3,19)

An index to a **memory**.

locator expression (3)

A sequence of **cells**, delimited on the left and right by a **locator**.

left locator (10)

The **locator** on the left of a **locator expression**.

luff (45)

Cells yet to be passed by the **dedicated register** of a **structured store**.

memory (19)

A function from **locators** to values.

opcode (14)

The first element of a **command sequence**, if this is a constant.

outside (49)

A **semantic equivalence** or **translation** between two **plans** holds outside a set of **locators** if the property holds **within** the complement of that set.

parallel block symbol (open) (87)

‘(|)’: indicates the start of a **parallel block**.

parallel block symbol (close) (87)

‘(|)’: indicates the end of a **parallel block**.

parallel block (87)

A set of **update schemes**. All **applicable update schemes** are applied simultaneously. If the resulting right hand side is not **consistent** none of the **schemes** are applied.

parameter resolution (60,62)

The mechanism by which parameters of an **archetype** are rewritten to evaluable expressions.

pointer (44)

The contents of **dedicated register**.

probation phase (101,107)

In an implementation, the phase in which the **guard** is checked.

programme counter (6)

The **register PC**.

programme state (46)

The **configuration** of a **programme store** and of the **register PC**.

programme store (46)

A **store** having the **programme counter** as **dedicated register**.

programme (46)

The contents of a **programme store**.

queue (45)

A **semi-structured store** which reads and writes in the same direction.

read access (44)

An **access** on the left hand side of an **update scheme**.

read dedicated (44)

A **dedicated register** which addresses only **read accesses**.

read structured store (44)

A **store** in which all **read accesses** are **structured**.

read to the left (44)

A **direct access** the **left locator** of which is invariant.

read to the right (44)

A **direct access** the **right locator** of which is invariant.

register (10)

A constant **locator**.

repeat (12)

“”: indicates a repeat of the previous left hand side.

resolution phase (106)

In an implementation, the phase in which values are derived for terms.

restoration phase (106,107)

In an implementation, the phase in which provisional instantiations are undone.

right locator (10)

The **locator** on the right of a **locator expression**.

satisfied (20)

A **configuration** is satisfied by any **configuration** of which it is a subset.

script (4)

An **initial configuration** and an **update plan**.

selection phase (106,108)

In an implementation, the phase in which a choice is made between proceeding to the **verification phase** or continuing with the **expansion phase**.

semantic equivalence (47)

A property holding between two **update plans**.

semantic irrelevance (50)

A property of **configurations** with respect to an **update plan**.

semi-constant (9)

A regular expression over constants.

semi-ground term (42)

A term for which a variable free value can be derived.

semi-structured store (45)

A store which is **read structured** and **write structured**, but not **fully structured**.

statically semi-ground (42)

A **semi-ground term** for which a variable free value can be derived without reference to the current **configuration**.

store structure (41)

A regular expression over store names.

store (3)

A two way countably infinite set of **cells**.

textual expansion (60)

The replacement text of an **archetype** call, before **parameter resolution**.

structured access (44)

An **access** exhibiting a certain structure.

structured store (44)

A **store** in which all **accesses** are **structured**, and in which all these **accesses** share a common **register**.

text (17)

A sequence of terms.

textual expansion (60)

The replacement of an **archetype** by its definition.

translation (47)

A property holding between two **update plans**.

type alias (40)

A type declaration.

type grammar (41)

A context free grammar based on the structure of types.

type primitive (40)

A type name not appearing as the left hand side of a **type alias**.

unprotected variable (110)

A term in a **converse** which is not **(semi-) ground**.

update plan (4)

A set of **update schemes**.

update rule (3)

An **update scheme** containing no variables.

update script (4)

An **update plan** and an **initial configuration**.

update scheme (3)

Consists of a left hand side, a right hand side (both **configurations**), and a **guard**.

verification phase (101,108)

In an implementation, the phase in which the **guard** is checked.

within (49)

A **semantic equivalence** or **translation** holds within a set of **cells** if the property holds when cells not in the set are ignored.

write access (44)

An **access** on the right hand side.

write dedicated (44)

A **dedicated register** which addresses only **write accesses**.

write structured store (44)

A **store** in which all **write accesses** are **structured**.

write to the left (44)

A **direct access** the **right locator** of which is invariant.

write to the right (44)

A **direct access** the **left locator** of which is invariant.

- λ
 - calculus, 7, 25–28, 116
 - MMachine, 25, 26, 28
- λ^+
 - calculus, 25, 32
 - MMachine, 25
- \parallel , 87, 88
- (\parallel) , 87, 88
- $\parallel)$, 87
- access, 44, 45
 - direct, 44
 - read, 44, 45
 - structured, 44
 - write, 44, 93
- addressing modes, 2, 7, 55, 56, 72, 77, 78, 81, 85
- alternatives, 15, 16, 98, 101, 109
- ALU, 68
- ambiguity, 39
- applicable, 4, 6, 16, 21–23, 87–89, 95, 98, 101, 105, 108–110, 112, 117
- application
 - archetype, 60, 116
 - parallel block, 87
 - update rule, 4
 - update scheme, 21–23, 98, 116
 - update script, 28
- archetype, 5–7, 43, 46, 55–61, 63, 65, 67, 69, 72–79, 83, 89, 91, 92, 94–99, 102–105, 112, 115–117, 119, 121
- call, 43, 55–62, 65, 66, 93, 104–106
- definition, 56–62, 64, 66, 67, 96, 97, 102, 105, 106, 112
- expansion, 56, 58–66, 71, 74–79, 83, 89, 91, 92, 96, 98, 102–108, 112
- grammar, 41, 65, 66, 79
- index, 57, 58, 60
- parameters, 43, 55, 56, 58–60, 62–66, 102–106, 108
- ambidextrous, 59, 60, 97
- ambiguous, 111, 112
- command, 57, 60, 112, 113
- deterministic, 108
- left-handed, 56, 59
- nondeterministic, 108
- recursive, 56, 63–65, 103
- right handed, 59
- arithmetic
 - instructions, 72–75, 78, 90–92
 - operators, 43, 73, 91

- attribute grammars, 55
- background storage device, 46
- backtrack rule, 111
- backtracking, 39, 95, 98, 108–112
- base value, 72, 77, 81–84
- bounded nondeterminism, 41
- box diagrams, 48–52, 82
- C, 37, 38, 43
- canonical form, 95–97, 101–103, 106
- carry bit, 92
- cast, 40, 41
- cell, 3, 10, 19, 20, 22, 42, 45, 49, 51, 52, 64–66, 79, 82, 83, 85, 96, 101, 109, 111, 112
 - lee, 45, 52
 - luff, 45, 52
- close parallel block symbol, 87
- closure, 9, 23, 41, 67
- code generation, 2, 71, 94
- coercion, 40
- combinatory logic, 116
- command, 6, 8, 12–16, 46, 67, 102, 112
 - driven, 14
 - sequences, 14
 - set, 55
- commands, PDP-11, 4
- comparison operators, 72
- compilers, 48, 71, 116, 117, 119, 121
 - C, 38
 - FLIP, 25–28, 36, 37
 - Update Plan, 37
- completed left hand sides, 109
- concatenation operator, 43
- configuration, 3, 4, 14, 19–23, 26, 29, 33, 35, 37, 42, 43, 47–52, 54, 56, 63–65, 80–83, 88, 89, 93, 97, 101, 103, 107, 109, 110, 116
 - final, 4
 - initial, 4, 10, 13, 16, 23, 26, 28, 29, 85, 93, 100, 106, 107
- consistency, 3, 4, 20, 22, 89, 95, 101
- consistent substitution, 62, 96
- constructor, 28, 32–36
- context, 59–61, 83
- context free grammar, 10, 13, 17, 41, 56, 64, 65, 88
- context free languages, 64
- converse, 109–112
- correspond, 110
- data transfer commands, 72, 78
- derivation, 4, 23, 49–52
- derived attribute, 55
- development, 23
- diagram chasing, 50, 52, 53, 82
- don't care, 60, 97, 103, 107
- dummy, 78
- equivalence, 47–51, 53, 72, 84
 - semantic, 36, 47–49, 72, 84–86, 94
- expansion
 - phase, 105, 106, 108
 - archetype, 60, 65, 66
 - textual, 58, 60, 61, 63, 65, 74, 75
- Extended Affix Grammars, 69
- final
 - configuration, 4
 - development, 4, 11, 13, 16, 23, 28, 50
- finite
 - ambiguity, 116
 - semi-constant, 9
- FLIP
 - intermediate code command, 30, 36
 - machine command, 27, 28, 30–33, 36, 37
 - programme, 26, 34, 35
- grammar driven recognition, 64
- graph reduction, 25
- graphical display device, 46
- ground term, 21, 22, 42, 43, 76, 96, 110
- grounding, 11, 21, 43, 46, 63, 66, 67, 69, 76, 105, 106, 110, 112

- analysis, 96, 99, 105
- guard, 3, 4, 12, 15, 16, 21, 22, 43, 45, 56, 59–64, 88, 95, 96, 100, 101, 104, 107, 108
- heap, 28, 31, 39, 45, 52, 99, 103, 106
- implementation, 7, 39, 41, 43, 66, 87, 95–99, 102–105, 108, 109, 111, 112, 115–117, 120, 122
- inconsistency, 20, 87, 101
- indirection, 16, 81, 97, 99, 103, 108
- inherited attribute, 55
- initial configuration, 4, 10, 13, 16, 23, 26, 28, 29, 85, 93, 100, 106, 107
- initialisation, 26, 28, 93
- input, 12, 13, 28
 - stream, 12, 13, 28, 37, 39, 45
- instantiation, 3, 4, 6, 15, 21–23, 41–43, 64, 87–89, 91, 99, 102, 103, 105–108, 112, 113
- instruction
 - cycle, 68, 92, 100, 106
 - register, 68
 - set, 2, 68
 - FLIP, 26–29
- intermediate code, 26–32, 35, 38, 71, 78, 79, 119
- internal
 - forwarding, 90, 94
 - representation, 31, 95, 97–101, 104, 108
- irrelevance, 49–52, 54, 82
 - semantic, 49–52
- ISPS, 4, 5
- label, 78, 81, 84, 85
- lee, 45, 52
- left hand side, 3, 4, 12, 14, 15, 19, 21–23, 34, 42–45, 48, 58–61, 65–67, 88, 89, 93, 102, 104–106, 108–111
 - context, 58, 61
 - expansion, 60
- archetype, 60, 65, 66
- left locator, 10, 11, 61, 68, 91, 96, 97, 99, 101, 102, 104, 107, 110, 112
- left-handed archetype, 56, 59
- length operator, 43
- locator, 3, 6, 10–14, 19, 26, 36, 37, 39, 40, 44, 51, 53, 56, 59, 66, 72, 77, 79, 81, 82, 84, 97
 - expression, 3, 6, 10, 19–22, 30, 34, 42, 44, 58, 59, 61, 76, 83, 93, 99, 101, 104, 105, 110
 - store, 99
 - left, 10, 11, 61, 66, 68, 91, 96, 97, 99, 101, 102, 104, 107, 110, 112
 - right, 10, 11, 15, 61, 66, 96, 99, 104, 110
- luff, 45, 52
- macro, 56, 89, 103
- mappings, 43
- memory, 19
 - access, 97
 - address register, 68
- nfib**, 37, 38
 - number, 37, 38
- non-ground term, 14, 42, 63, 74, 110
- nondeterminism, unbounded, 41, 42
- opcode, 7, 14, 16, 91
- open parallel block symbol, 87
- output, 12, 13, 32, 51
 - stream, 13, 36, 37, 39, 45, 51
- outside, 49, 51–53, 82
- parallel
 - block, 87–89, 93, 95, 98, 108, 112, 115, 117
 - close symbol, 88
 - open symbol, 88
 - processors, 55, 119–122
 - asynchronous, 7, 68, 69, 87, 115, 117

- synchronous, 7, 69, 87, 115
- parallelism, 94, 116, 117
- parameter
 - resolution, 62–65
 - stage, 60
 - store, 104
- PDP-11, 2, 4, 68, 71
- peephole optimiser, 5, 85
- pipeline symbol, 87, 88
- pipelining, 7, 87, 90, 92–94, 120, 122
- pointer, 12, 13, 28, 33, 34, 44, 45, 50, 54–56, 90, 95–97, 99, 100, 103, 104, 106, 107
 - read, 45
 - write, 45, 52
- programme, 1, 26, 46, 115, 116, 119
 - counter, 6, 13, 14, 34, 67, 68, 112
 - equivalence, 1, 47, 94, 116
 - flow commands, 72, 74, 84
 - state, 46
 - store, 26, 39, 46, 54, 82, 84, 85
 - transformations, 54
 - FLIP, 26, 34, 35
- programming, structured, 115
- queue, 45, 52
- read
 - access, 44, 45
 - dedicated, 44
 - only store, 46, 52
 - pointer, 45
 - structured store, 44, 45
 - to the left, 44
 - to the right, 44
- recursive archetype, 63, 103
- register, 10, 44
 - allocation, 7, 71, 81
 - dedicated, 44, 45
- regular expression, 9, 40, 41
- repeat, 12
- resolution, 74
 - phase, 106
- restoration phase, 106
- rewrite systems, 2
- right hand side, 22
- right handed archetype, 59
- right locator, 10, 11, 15, 61, 66, 96, 99, 104, 110
- RISC II, 7, 87, 90, 92
- satisfying, 13, 16, 20, 22, 28, 83, 97, 106, 107
- scheme counter, 98
- seed, 49, 50
- selection phase, 106, 108
- self-modifying code, 46
- semantic
 - equivalence, 36, 47–49, 72, 84–86, 94
 - irrelevance, 49–52
 - properties, 7, 39, 44, 46–48, 54, 71, 115, 119
- semi-constant, 9
- semi-ground term, 21, 42, 43, 63, 66, 96, 108, 110–112
- semi-structured store, 45
- stack, 12, 28, 29, 31–34, 39, 44, 45, 50, 52, 54, 71, 73, 84, 85, 111
 - archetype, 108
 - backtrack, 109, 111, 112
 - compile time, 26, 29, 32, 33
 - tripwire, 34, 35
- standard environment, 73, 97, 100, 106, 112
- static store, 28, 39, 45
- store, 3, 19, 28, 48, 51, 54, 79
 - declaration, 41
 - name, 41
 - structure, 41
 - locator expression, 99
 - parameter, 104
 - programme, 26, 46, 81, 82, 84, 85
 - read-only, 46, 52
 - static, 28
 - structured, 44, 45
 - read, 44, 45
 - write 44, 45, 52
 - value, 95, 96, 99, 103
 - write-only, 46, 52
- structured

- access, 44
- programming, 115
- store, 44, 45
- substitution, 3, 21–23, 63, 74–76, 88
 - consistent, 62, 96
- suspension, 28–35
- syntactic sugar, 2, 6, 7, 9, 11, 14, 15, 17, 25, 56, 57, 59, 60, 67, 95, 96
- syntax macros, 56
- synthesised attribute, 55
- textual
 - expansion, 58, 61, 63, 65, 74, 75
 - stage, 60
- translation, 47–54, 71, 79, 82, 85, 93
- type
 - alias, 40, 41
 - grammar, 41, 43, 65
 - name, 41
 - primitive, 40
- typing, 6, 7, 10, 31, 39–46, 64, 65, 85, 97, 119, 121
- unbounded nondeterminism, 41, 42
- update, 19, 22, 48, 49, 82, 83, 88, 89, 109, 110
 - plan, 2–5, 6, 9, 14–16, 19, 23, 25, 28, 38–42, 44, 46–49, 51, 54, 60, 64, 65, 67, 85, 87, 89, 95, 98, 102, 105, 106, 108, 109, 111, 115, 117
 - rule, 3, 4, 6, 41, 64, 87, 89, 109, 116
 - scheme, 3, 4, 6, 8, 10, 11, 13, 14, 16, 19–23, 26, 27, 32, 35, 37, 38, 40–45, 55, 56, 58–62, 64, 65, 67, 74–80, 83, 84, 87–89, 91–93, 95–102, 105–112, 115–117
 - script, 4, 10, 16, 23, 26, 28, 30, 37, 47, 48, 54, 85, 89
- Update Plans, 2–5, 7–11, 17, 19, 25, 37, 39, 41, 55, 69, 71, 85, 87–90, 95, 97, 98, 108, 115–117, 119–122
- semantics, 2, 4, 7, 9, 19, 88, 89, 115, 116, 119
- syntax, 2, 4, 7, 9–11, 17, 41, 56, 57, 88, 115, 119, 121
- value store, 95, 96, 99, 103
- whitespace, 9, 10, 87
- within, 49
- write
 - access, 44, 45, 93
 - dedicated, 44
 - only store, 46, 52
 - pointer, 45, 52
 - structured store, 44, 45, 52
 - to the left, 44