# PDF hosted at the Radboud Repository of the Radboud University Nijmegen 

The following full text is a publisher's version.

For additional information about this publication click this link.
http://hdl.handle.net/2066/113882

Please be advised that this information was generated on 2021-07-16 and may be subject to change.

## FUNCTIONAL PROGRAMS

## AS

## EXECUTABLE SPECIFICATIONS



Pieter Koopman

## FUNCTIONAL PROGRAMS

## AS

## EXECUTABLE SPECIFICATIONS

een wetenschappelijke proeve op het gebied van de wiskunde en informatica.

## Proefschrift

> ter verkrijging van de graad van doctor aan
> de Katholieke Universiteit te Nijmegen, volgens besluit van het college van decanen in het openbaar te verdedigen op maandag 10 december 1990 des namiddags te 3.30 ur

door

## Petrus Wilhelmus Maria Koopman

geboren op 3 juli 1957 te Nijmegen

Druk: Krips Repro Meppel

Promotor: Prof. dr. H.P. Barendregt
Co-promotor: Dr. ir. M.J. Plasmeijer
(c) 1990, Pieter Koopman, The Netherlands. ISBN 90-9003689-X

## Manuscriptcommissie:

Prof. dr. C.C.A.M. Gielen Katholieke Universiteit Nijmegen
Dr. ir. S.M.M. Joosten
Dr. W.G. Vree Techinische Universiteit Twente Universiteit van Amsterdam

Graag wil ik hierbij iedereen bedanken die op een of andere manier heeft bijgedragen aan de totstandkoming van dit proefschrift. Zonder jullie hulp zou dit proefschrift nooit geschreven zijn.

## Table of contents

1 Introduction ..... 7
1.1 Place of specifications in the implementation process ..... 8
1.2 Specification formalisms ..... 13
13 Specifications treated in this thesis ..... 15
2 Operational Machine Description ..... 19
21 The description method ..... 20
The machine state ..... 20
The micro-instructions ..... 21
The instructions ..... 22
Program execution ..... 24
22 Description of a sımple machıne ..... 25
The state of Mac-1 ..... 27
The micro-instruction level of Mac-1 ..... 27
The Mac-1 instructions ..... 29
Execution of Mac-1 programs ..... 30
An example Mac-1 program ..... 31
23 Discussion ..... 32
3 The ABC-Machine ..... 35
31 The architecture of the ABC-machine ..... 37
32 The micro-instruction set ..... 38
The graph store ..... 39
The descriptor store ..... 40
The A-stack ..... 41
The B-stack ..... 41
The program ..... 42
The program counter ..... 43
The C-stack ..... 43
The input-output channel ..... 43
33 The instruction set ..... 44
Graph manipulation ..... 44
Manipulating the A-stack ..... 46
Manipulating the B -stack ..... 47
Changing the flow of control ..... 47
Generating output ..... 48
3.4 Program execution ..... 49
The instruction cycle ..... 49
Booting the machine ..... 49
3.5 ABC-assembler ..... 50
3.6 Comparison with the G-machine ..... 51
The G-machine state ..... 52
The G-machine instructions ..... 53
Initial and final state of the G-machine ..... 54
Differences between the abstract machines ..... 54
Differences between the descriptions ..... 55
3.7 Discussion ..... 55
4 Abstract Program Translation ..... 57
4.1 The description method ..... 58
4.2 Bracket abstraction ..... 60
4.3 The traditional description of bracket abstraction ..... 61
The basic abstraction algorithm ..... 61
Improving the generated combinator code ..... 63
Data structures ..... 65
Local definitions ..... 65
Recursion ..... 66
Summary of the abstraction algorithm ..... 67
Suitability of the description method ..... 68
4.4 Bracket abstraction described by syntactical matching ..... 68
Combinators used in the abstraction algorithm ..... 69
The expressions handled ..... 69
The abstraction algorithm ..... 70
Improving the generated combinator code ..... 71
Reduction ..... 71
Suitability of the description method ..... 72
4.5 Bracket abstraction as abstract program transformation ..... 73
The expressions handled ..... 73
Combinators used in the abstraction algorithm ..... 73
The abstraction algorithm ..... 74
Improving the generated combinator code ..... 75
Executing the generated combinator code ..... 75
Suitability of the description method ..... 76
4.6 Discussion ..... 77
5 Translating Clean to ABC-Code ..... 79
5.1 Graph rewriting in Clean ..... 80
5.2 Graph rewriting on the ABC-machine ..... 82
Construction of the stack frame ..... 83
Rule alternative entries ..... 84
Example of the rewriting process ..... 87
5.3 The AST for Clean ..... 89
5.4 Compilation ..... 90
Descriptors ..... 91
Preparing the arguments ..... 92
Rewrite alternative entries ..... 93
The additional rewrite alternative entry ..... 99
5.5 The run-time system ..... 100
5.6 Curried functions ..... 101
5.7 Optimizations ..... 103
The extended compilation scheme using the B-stack ..... 103
Other optimizations ..... 106
Efficiency gain of the optimizations ..... 107
5.8 Discussion ..... 109
6 Artificial Neural Networks ..... 111
6.1 Description tools for artificial neural networks ..... 113
The values processed ..... 113
Neuron characteristics ..... 114
Neuron interconnection patterns and evaluation orders ..... 117
Learning rules ..... 120
Description of specific artificial neural networks ..... 124
6.2 The Adaptive Linear Combiner ..... 125
6.3 The Adaline ..... 125
6.4 The Perceptron ..... 127
The single layer perceptron ..... 127
The multi-layer perceptron for continuous values ..... 129
The multi-layer perceptron for binary values ..... 131
6.5 Learning Vector Quantization ..... 134
6.6 Feedback Networks ..... 138
Convergence ..... 139
Avoiding local minima ..... 142
6.7 Discussion ..... 145
7 Summary and Conclusion ..... 147
A Definition of the ABC-Machine ..... 151
A. 1 The micro-instruction set ..... 151
A. 2 The instruction set ..... 159
A. 3 ABC-assembler ..... 168
B Translating Clean to ABC-code ..... 175
B. 1 Reduction entries ..... 177
B. 2 Code for the left handside of a rule alternative ..... 178
B. 3 Code for the right handside of a rule alternative ..... 179
B. 4 Auxiliary functions ..... 182
C Artificial Neural Networks ..... 185
C. 1 Network description tools ..... 185
C. 2 A feedback network to solve the N -queens problem ..... 191
References ..... 193
Index ..... 201
Samenvatting ..... 205

## Chapter 1 Introduction

1.1 Place of specifications in the implementation process
1.2 Specification formalisms
1.3 Specifications treated in this thesis

In 1978 Backus advocated the use of functional programming languages as the new programming discipline needed to solve the problems in software production [Backus 78]. This initiated a large amount of research to achieve at least a reasonable execution speed for these languages. The scope of this research ranges from theoretical studies on the computational models used in the functional languages to the survey and construction of special purpose hardware to execute functional languages in a sequential or parallel way. In the mean time a number of nice languages has been developed. The actual use of functional programming languages received much less attention than the implementation and the study of theoretical properties, partly caused by the bad performance of implementations available.

In this dissertation the suitability of functional programming languages as a vehicle for specifications is investigated. Functional languages seem to be suited as executable specifications due to the high abstraction level and powerful language constructs available. Such a programming language has as advantages over a tailor made specification formalism the well defined semantics and the possibility to check the specifications on consistency by a compiler. Furthermore the specifications are executable. So, the specifications are their own prototype implementation. The suitability of functional languages as specification formalism is illustrated by using them in daily practice on the department of computational models and parallel systems of the Nijmegen university. Sequential and parallel implementations of functional languages are developed here. The underlying computational models of these languages as well as other computational models, like neural networks, are investigated.

In this introductory chapter a motivation of the work presented in this thesis is given. First the function of specifications and prototypes in the construction of software is clarified in section 1.1. Then we argue why functional programming languages are very well suited to be used as a formalism for specifications. Finally, an overview of the specifications presented in this dissertation is given.

### 1.1 Place of specifications in the implementation process

To clarify the role and importance of specifications in an implementation process, a division in logical phases is made. The implementation process is only one step in a software engineering process. Usually, it is preceded by a study on the kind of product needed and it is accompanied by training and documentation upon its introduction. Afterwards the product still needs maintenance for some time.

The production of an implementation, in hard or software, can be split in a number of phases:

1. Determination of the requirements; specification of the problem to be solved.
2. Design of an algorithm, that meets these requirements.
3. Construction of the product. An implementation of the specified algorithm is made.
4. Verification of the product. This phase consists of testing (or proving) whether the product meets the requirements or not. Errors made during the construction are spotted now.
In an ideal situation these phases are strict sequential and the interfaces are unambiguous. But, in many situations these phases are not clearly separable and sequential. The definition of the requirements depends on what is feasible during the implementation and is adjusted as soon as the first tests behave otherwise as expected. Also the interface between the phases is usual not free of misapprehensions. Nevertheless, these phases reflect the four main points in an implementation project; what must be solved, how must it be solved, construction of the product, and its verification. In large projects this division in phases can be used for sub-tasks as well, this makes the phase to which an activity belongs depending on one point of view.

Firm specifications are needed to construct an unambiguous interface between the phases. The interface between the construction and verification phase is usually the best defined one. By definition an implementation in a computer language is a formal specification, the program can be checked syntactically and and its behaviour can be studied. In order to prevent misunderstandings the use of formal specifications for the other interfaces is recommended. The specifications occurring in this process are: specification of the requirements, specification of the algorithm and specification of its implementation.

Analysis shows that the verification phase often requires very much time. Not only errors made during the construction of the final product are encountered, but also wrong design decisions or erroneous requirements can be discovered now. In the literature alarming figures are reported; more than half of the errors originates from wrong requirements and over a quarter of the errors in the product are due to the design [Beizer 83, DeMarco 78]. It is usually very expensive to cope with the latter ones. A prototype can be built to prevent this late detection of wrong requirements or an erroneous design.

## ad 1: Determination of the requlrements

Ideally the functionality of the product is defined in an unambiguous formalism. An ordinary mathematical specification is often used, but such a specification is not always applicable. A number of special purpose formalisms is developed to be used in those situations. In other situations it is not possible to give a precise definition at all. In those situations the requirements must be informally defined, e.g. an abstract machine to express the implementation of graph reduction conveniently.

Sorting items in ascending order is an example where the language of mathematics is very well suited to describe the intended property of the sequence of items, but less suited to describe how such a sorted sequence can be obtained. Provided that proper comparison relations, $\leq$, are defined on the items of a sequence and their indices one can define:

A sequence $S$ of items $I_{l}, I_{2}, \ldots, I_{n}$ is sorted iff $\forall i, j \in\{1,2, \ldots, \mathrm{n}\}: I_{i} \leq I_{j} \Leftrightarrow i \leq j$.
This specifies fully what is meant by a sorted sequence, but not how it can be constructed from an arbitrary sequence. As far as this definition is concerned a perfect way to obtain the sorted sequence is to generate all possible sequences and pick the first one that obeys the requested property. A constructive specification, an algorithm, is needed to show how a sorted sequence is obtained in a more efficient way. The static specification of the desired properties can be used to prove that the algorithm is correct.

In many situations a formal specification of the required properties cannot be provided, e.g. in general there is not a formal specification of the required properties of a (abstract) machine available. In other situations the proof of the correctness of the algorithm is so elaborated that it is omitted in every-day practice. Given an accurate description of the programming languages $A$ and $B$ and a compilation scheme from $A$ to $B$, it is in principle possible to prove its correctness, but in practice this proof is usually omitted due to its length and complexity.

The design of the algorithm is described by a specification. Experience ${ }^{1}$ shows that ambiguities can only be avoided by a formal specification. The design can be the result of an evolution of a sequence algorithms with an increasing algorithmic efficiency.

A specification of the design is a concise description of a constructive algorithm with the following properties:

- Consistent and correct, otherwise it is hard to expect a satisfying product.
- Sufficiently complete, the specification must contain sufficient information to construct the actual implementation, but matters like execution speed or space requirements are usually not fully determined.
- Unambiguous, any doubts about the exact meaning of the specification severely reduce its value, therefore formal languages are required as carrier of the specification.
Using a suited formalism the construction of the formal specification of the algorithm can be helpful in the construction of its design.

When a formal definition of the requirements is available the specification of the algorithm can be proven to be correct, or the correctness of the specification can be assured by deriving the program from the requirements. The length of the proofs and derivations limit the use of these verification techniques to small programs.

A well known constructive algorithm to sort a sequence of items is quicksort [Hoare 62]. This algorithm splits a sequence of items into two parts which are then sorted separately. To split a sequence an element, called median, of the sequence is chosen. The sequence is split into a fraction of elements less or equal to the median and a fraction of elements greater than the median. These fractions are sorted separately. The sorted sequence is obtained by the combination of both sorted fractions. The algorithm terminates by the observation that an empty sequence is always sorted.

$$
\begin{aligned}
\text { quicksort } \varnothing= & \varnothing \\
\text { quicksort } S= & \text { quicksort }\left\{e \mid e \in S^{\prime} \wedge e \leq m\right\} \cup\{m\} \cup \text { quicksort }\left\{e \mid e \in S^{\prime} \wedge e>m\right\}, \\
& \text { where } S^{\prime}=S \backslash m \text { (the sequence } S \text { without element } m \text { ), for some } m \in S S
\end{aligned}
$$

Here the set notation is used to denote sequences, although the order of elements is relevant here and S might be a multi-set. This specification ${ }^{2}$ defines the actions to be performed, but not their order nor the way the elements must be stored.

[^0]When there is any doubt about the determination of the requirements or the specified algorithm, it makes sense to make one or more prototypes. Prototype implementations can be used to study the dynamic behaviour of the specified system. In this way it can be investigated whether the system will act as required or not.

## Prototypes

A prototype is a (partial) implementation to test some key aspects of the intended, probably complex, algorithm. It obeys some aspects of the general specification. The differences between the final product and the prototype implementation are the low efficiency of the prototype, limitations in the amount of data that can be processed and a limited set of the required features. Sometimes a prototype contains additional capabilities to test the specification or to perform some measurements. In order to be useful as prototype it must at least be able to execute small scale tests at a reasonable speed.

The number of prototypes needed depends on the problem. When the requirements and the way to fulfil them are absolutely clear no prototypes are needed. In other cases one or more prototypes are constructed, each one showing some aspects of the product.

Prototypes must be constructed in a high level language for two reasons. First of all the implementation language must be close to the specification language in order to avoid errors in this implementation. When the prototype does not behave as intended it is not clear whether to blame the specification or its implementation. Secondly, the effort needed to construct a prototype must be significantly lower than the amount of work involved in the actual implementation, otherwise it is better to omit the construction of a prototype.

A Miranda ${ }^{3}$ prototype of the quicksort algorithm defined above is shown (the sequences are represented by lists and the first element of such a sequence is used as median):

```
quicksort :: ["] -> ["]
quicksort [] \(=\) []
quicksort (m:s) = quicksort \([e \mid e<-s ; e<=m]++[m]++\) quicksort \([e \mid e<-s ; e>m]\)
```

When this algorithm is compared with the specification above only a suited data type to represent sequences and a choice for the median are additionally specified. This data structure happens to be better suited to denote sequences than the set notation. Memory use and order of actions is not specified, only the computations needed are indicated.

[^1]After the design of the algorithm and its validation by the prototype, the implementation phase starts. Although a prototype can be a complete implementation it is only in rare cases useful as product. The product must obey several constraints on run time, space consumption and the amount of data to be processed. Also appropriate interfaces to the user and the system must be constructed. Among the topics to be solved are: efficient implementation of the used data structures, choice of the flow of control, efficient implementation of the specified algorithm, modular structure and documentation.

For a product in software also the programming language must be determined. Some of the criteria to chose an implementation language are: the availability of an appropriate implementation; portability; familiarity; suitability for the designed algorithm and speed of execution. When speed is an important aspect of the final product the actual implementation is usually made in an imperative language. Using currently available compiler technology it is hard or impossible to beat imperative languages on most stock hardware. When the speed constraints are less severe and/or more efficient compilers are available other kind of languages can be used for the implementation.

In our sorting example the space requirements and the number of sweeps through the sequence can be improved, also a better choice for the median is feasible. It is possible to specify these matters formally, but usually this is not needed nor wanted. These implementation issues are not essential for the quicksort algorithm, only for a fast special implementation of it.

The product corresponding to quicksort specification depends on the number and the kind of items to be sorted. For a small amount of data the Miranda prototype defined above can directly be used as the final product. But, often speed and space use do matter. Then a real product obeying the specification has to be encoded. An efficient imperative algorithm written in Pascal [Wirth 71] is given below. It sorts the global array a in situ, elements in the array are reshuffled to obtain sets of elements greater and less than the median in an efficient but rather hard to understand way. These sets are recursively sorted.

```
procedure quicksort (l,r: integer);
var m, i j, tmp: integer;
begin
    If \(r>l\) then
    begin
\begin{tabular}{ll}
\(m:=a \mid r] ;\) & \{ median \\
\(i:=\vdash-1 ;\) & \{start for left part; elements less or equal to median \\
\(j:=r ;\) & \(\{\) start for right part; elements greater or equal to median \\
repeat &
\end{tabular}
        repeat
            repeat \(i:=i+1\) until \(a[i]>=m ; \quad\) \{ expand left part with well placed elements \}
            repeat \(j:=j-1\) until \(a[j]<=m ; \quad\) \{ expand right part with well placed elements \}
            tmp :=a[i];a[i]:=a[j];a[j]:=tmp; \{ swap elements to right part
        until \(j<=i\); \(\quad\{\) totally split
        \(\mathrm{a}[\mathrm{j}]:=\mathrm{a}[\mathrm{i}] ; \mathrm{a}[\mathrm{i}]:=\mathrm{a}[\mathrm{r}] ; \mathrm{a}[\mathrm{r}]:=\mathrm{tmp} ; \quad\{\) median to centre \}
        quicksort(l,i-1);
        quicksort(i+1,r);
    end
end;
```

In this product the memory use and the order of actions are exactly specified. It is optimized for space use and speed; sometimes some superfluous actions are performed in order to avoid a test in all situations.

## ad 4: Verification

The correctness of a specification is a serious problem. When some other formal specification is available a correctness proof can be given. For instance a sorting algorithm can be proven to be correct by showing that none of the elements is lost and that the final order of the elements obeys the sorting criterion. When no other specification nor formal requirements of the product are available a correctness proof cannot be given. Usually, correctness proofs appear to be much longer than the specifications, this limits the practical use of proofs to small programs. For large programs proving the correctness is too much work, moreover the value of a proof decreases when its length increases. For very large programs neither testing nor proving can give absolute certainty about the correctness. A proof needs to be correct and complete to prove anything, so the proof itself must be verified. The correctness of a proof is a serious problem, especially with very elaborated proofs. Proofs and testing can only increase the confidence in a product, but both can be very valuable [Joosten 89].

### 1.2 Specification formalisims

In order to be useful, a specification must be correct and understandable. In general, descriptions are better understandable when they are shorter. Descriptions can be made more compact by omitting irrelevant issues and by the
introduction of new description primitives. In order to keep the descriptions clear and easy accessible for many people the introduction of description tools must be done very reluctantly. The brevity of the descriptions is bounded by the expressive power of the description language. In order to be used for an unambiguous definition the specification language itself must be defined properly.

Many special purpose formal description systems are proposed and used nowadays. The advantages of these systems are obvious:

- by the introduction of special purpose notations very compact specifications are possible;
- since these systems are formal they are unambiguous.

However, there are also some disadvantages of the use of special purpose formal systems:

- in order to understand the specifications additional knowledge about the specification language is needed;
- mechanical checks on the syntax and consistency of the specified system are only possible after some implementation of the developed formalism;
- these systems are usually not meant nor suited for execution, so additional work must be done to construct a prototype.
The use of programming languages instead of a special purpose formalism for the specification has a number of attractive advantages.
- The syntactical correctness, absence of unbound identifiers and type consistency can be checked most efficiently and concisely by a compiler. This does not imply that a compiler is able to verify that the complete specification is correct, only partially correctness can be tested.
- Using an existing programming language there is no need nor place for new language constructs in the description. This implicates that the semantics of the description are clear and well defined.
- The specification is its own prototype implementation. This eliminates the effort to construct a prototype and to keep the specified product and prototype similar.
- The specification is checked by the corresponding prototype in a direct way, construction of a separate prototype is only an indirect test of the specification. Although each test shows just that the specified system behaves as shown on the test data, it yields valuable information on the dynamic properties of the system.
The obvious drawback of the use of most programming languages as specification language is the great implementation effort required. Much more time is spent on solving all kind of implementation details than on writing a specification of the product. The difference between the design of the algorithm and the construction of the product vanishes. Due to the implementation details specified in the same description as the algorithm it becomes too elaborate to handle.

Due to the high abstraction level and powerful language constructs functional languages seem to be suited better as specification formalism than imperative programming languages. The comparison of the quicksort algorithm and its prototype implementation in Miranda yields a first indication of the suit-
ability of functional languages as formal specifications. The advantages of using a functional instead of an imperative programming language for the specification are:

- Due to the powerful language constructs and high abstraction level the designer can spend (more of) his time writing the specification instead of solving implementation problems;
- Functional programs are much compacter than equivalent imperative programs, this significantly increases their value as specification.
- Using the abstraction mechanisms of functional languages a hierarchical description can be given. Layered descriptions enable detailed explanations together with a clear overall view.
- When a functional programming language with lazy semantics is used, only the data flow must be specified, the flow of control is deduced automatically from the data needed.
The difference between a functional description and an ordinary functional program is the emphasis on a clear explanation of the system in the description. Efficiency and algorithmic complexity of the description itself are irrelevant. A specification describes an efficient algorithm while an ordinary functional program itself must be efficient. These are just the usual differences between a prototype and the product.

In this thesis the suitability of functional programs as executable specifications is investigated. We are mainly interested in relatively large examples where correctness proofs have a limited applicability due to the size of the problems or to the absence of formal requirements to be met.

A general available functional programming language, Miranda [Turner 85], is used as specification language. Mainly the basic expression rewrite properties of this language are used. This subset is more or less common in all functional languages, so other languages could be used for descriptions as well. Although Miranda is not designed as a specification language, it is reasonably suited. In our opinion it is not worthwhile to design a new tailor made language for each application area, when an existing language is sufficiently fitted. Though Miranda is used with success as a description language, some general remarks concerning its design will be made. This critic concerns the lack of term rewriting semantics for partially parametrized functions. Unfortunately the formal semantics of Miranda has been announced for a long time, but it is still not published. Fortunately, the semantics of a large part of the language is simple. So, a formal specification of the semantics will be close to the language [Fehr 89]. Moreover, only one implementation of Miranda is allowed so it can be used as a, slow, reference implementation.

### 1.3 Specifications treated in this thesis

A large part of the specifications described in this thesis has actually been used to describe real products in a large implementation project of functional lan-
guages on several (parallel) architectures currently under execution at the University of Nijmegen.

As a first step in the implementation process, a functional program is transformed to a graph rewriting system. Graph rewriting systems are a well-defined computational model for functional programming languages [Klop 87, Eekelen 88]. The traditional computational model used for functional programming languages is the $\lambda$-calculus [Church 32, Barendregt 84], see for instance [Field 88] for a thorough overview. The advantages of the use of graph rewriting systems as computational model instead of the $\lambda$-calculus are: the notions of named (recursive) functions, pattern matching is included in the computational model and graphs enable the detailed description of the sharing of computations.

The second step in the compilation process is the transformation of the graph rewriting system to code for an imperative abstract graph rewriting machine: the ABC-machine. This enables reasoning about implementation issues not expressible in the functional graph rewriting model. The precise flow of control can only be described in an imperative model.


Fig 1.1 The compilation path for functional languages.
The translation of Miranda to Clean and the translation from Clean to ABC-code are specified Miranda. Also the abstract ABC-machine is described in this functional language. In this thesis the specification methods used are described and compared with other descriptions. To show that these description methods can be applied in practice and how comprehensible they are, the specification of the ABC-machine and the translation of Clean to ABC-code are treated.

In chapter 2 it is shown how functional programming languages can be used to describe an (abstract) machine. Such an executable specification is compared with a more common imperative description. The proposed description method is used in the next chapter for a thorough description of the ABCmachine.

A description method for translations is introduced and compared with other descriptions in chapter 4. This description method is employed in chapter 5 to specify the translation of the graph rewriting systems to imperative ABCcode. The translation of Miranda to Clean is described in a similar way
[Koopman 87, Plasmeijer 91], but boils down mainly to the removing of syntactical sugar and is hence less interesting than the translation described here.

Also other aspects of these translations, like type checking [Bakel 90], are described analogously. Currently the implementation of a parallel variant is implemented and described similarly [Smetsers 89, Plasmeijer 91].

Functional specifications can also be used outside the world of functional languages and their implementations. Chapter 6 contains a detailed description of a number of important artificial neural networks. These layered models are constructed of very simple processing elements. Usually these networks are described by a mathematical definition of the input-output relations of the various neurons. The dynamic adaptation of these networks can also adequately be described in a functional language.

Functional languages can be used as programming language or description language in many other fields [Hughes 89], e.g. in [Koopman 88] the use of interactive programs in a functional language is treated. Functional languages can also be used to describe many kinds of systems, for instance digital and analog circuits [Boute 88]. The use of functional programs as specification and prototype derived from a mathematical specification is treated in [Joosten 89].

The last chapter of this thesis contains our conclusions of the suitability of functional programming languages as description languages.

This thesis does not contain an introduction to functional programming. Nowadays there are a number of books containing such an introduction available, we recommend [Bird 88, Plasmeijer 91, or Field 88] for a thorough introduction. The first book is one of the few books treating functional languages without treating their implementation. The second book uses Miranda as functional language and shows that graph rewriting systems are a very suited computational model for the (parallel) implementation of functional languages. The last book uses Hope [Burstall 80] as functional language and uses the $\lambda$-calculus as computational model for the implementation. The language Clean is described only informally in chapter 5 of this thesis. A number of descriptions is available in the literature [Brus 87, Barendregt 87b, Eekelen 89, Plasmeijer 91].

# Chapter 2 Operational Machine Description 

### 2.1 The description method <br> 2.2 Description of a simple machine <br> 2.3 Discussion

In this chapter functional programming languages are used to describe (abstract) machine architectures. For each machine there are many possible descriptions, each one showing different aspects of the machine, e.g. the instruction repertoire, the micro-programming level or the hardware. The intention of the description presented in this chapter is to show the components and to explain their interactions as relevant from a programmers point of view.

The machine description consists of an implementation of the machine in a functional language. This implementation is not meant to be efficiently executable, but to supply a clear view of the relevant machine aspects. The description of machines in a functional programming language has a number of advantages over conventional descriptions. The description is an ordinary functional program. This implies that the description is executable and serves as a prototype machine implementation. The provided implementation appears to be valuable to study the dynamic behaviour of the machine. Using this prototype implementation the machine can be run in an early state.

Using the abstraction mechanisms of the functional programming language a hierarchical machine description can be given. Such a layered machine description enables a detailed explanation, without getting an overall view glutted by too much details at the top level of the description. The description can be chosen such that the machine aspects of prime interest are clearly shown, while other aspects are hidden. In this way one can abstract from the implementation of some machine components on one level, but accurately specify that object on another level. This hierarchical description method introduced in this chapter is generally applicable. It can also be used to describe large concrete machines, although they tend to have more parts than the machines shown in this thesis, hence they have a more elaborate state.

Due to the use of a functional programming language as description formalism, the descriptions are clearer and less error prone than conventional descriptions. The descriptions give valuable guide-lines to achieve an efficient implementation and serve as reference implementation of the described machine.

Commonly used description methods for (abstract) machines are:

- textual description, these specifications are hard to get unambiguous and correct;
- by the state changes in an imperative fashion, see this chapter for a comparison;
- by a definition of the semantics of the instruction set, see chapter 4 and 5 for a discussion of semantical descriptions.
After the introduction of the description method in section 2.1, a small but complete example is given to illustrate the description method proposed here. The example is taken from a textbook of computer architecture where it is used to illustrate the concept of micro-programming. Here only the instructions of the conventional machine level are specified to illustrate the proposed description method. The micro-programming level can be described similarly. In the next chapter this description method is used to specify an abstract graph rewrite machine as a large and more complex example.


### 2.1 The description method

The machine descriptions introduced in this chapter primarily define the semantics of the machine instructions. In general such a specification is not in close correspondence with possible hardware implementations of the machine described. This implies that the description usually does not reflect the common von Neumann computer architecture (a Central Processing Unit (CPU) and a memory interconnected by a bus), but shows the the machine components relevant from the programmers point of view. All machine components together determine the state of the machine. Instructions change the machine state. The machine descriptions introduced here, will be based on a specification of the state of the machine.

The instructions are described in a two level model. On the top level, the instruction level, the exact state transition of each instruction is specified in terms of micro-instructions. The micro-instructions form the bottom level of the description and define primitive operations on the logical components of the machine.

## The machine state

The machine description method presented in this chapter is applicable to sequential, imperative machines. Such a machine consists of a number of constructing parts (like various pieces of memory and input/output channels). The
state of the machine is completely described by the contents of the constructing parts. The state and hence the contents of these parts can be manipulated by the instructions defined on the machine.

For the description, the complete state of the machine is recorded in a data structure, called state. For every component of the machine a suited (abstract) data type is used, to describe its status or contents.

For instance, the state of a machine with one register (reg), a program counter ( $p \mathrm{c}$ ), a stack (stack) and a memory (mem) is described by:
(pc, reg, stack, mem)
The definition of the state affects the architecture as it is presented to the user. For example, when it is important to show that the stack is maintained with a stack pointer (sp) in the global memory (mem), the used machine architecture must be slightly different. This is reflected in the state of this machine, it is described by:
(pc, reg, sp, mem)

These two machine models can be depicted as:


Fig 2.1 The machine with a stack, a program counter, one register and a memory is depicted on the left. On the right hand side a similar machine is drawn, but here the stack is maintained with a special register containing a stack pointer in the memory.

It is not possible to use the state corresponding to the left machine to describe the right architecture. Adjustable data structures, like memories and stacks, must always be constructed in a tree-like structure. This shows clearly the dependencies between the machine components, which is generally desirable.

## The micro-instructions

To access the data structures describing the components of the machine tailor made functions are used. These functions are called micro-instructions, since they will form the building blocks of machine instructions. These micro-
instructions define an abstract data type describing the machine component at hand, e.g. a stack can be described by the functions to push and pop elements and to access an element on the stack. Only in extremely simple cases the use of the abstract data types can be omitted, e.g. registers with ordinary integer arithmetic can sometimes be modelled by the language implementation of integers.

A set of micro-instructions defining the stack for the first example machine architecture above is:

```
abstype stack
    with s_lop := stack -> value
        s_pop = stack -> stack
        s_push : value -> stack -> stack
```

Micro-instruction names are usually prefixed with some characters indicating the data structure manipulated.

A set of micro-instructions defining the memory is:

```
abstype mem
    with get_num :: address -> mem -> num
        get_ins = address -> mem >> mstruction
        m_update : address -> num -> mem -> mem
```

To create an executable specification all machine components have to be implemented. The implementation of the data type can be kept simple when the structure of the associated machine part is not a topic of interest in the given description, e.g. a stack can be represented as a list. The data type can also show a more accurate model of the machine component, e.g. the stack can also be represented by an array of values and an index in that array to represent the stack pointer.

Other micro instructions used in this introduction have similar definitions. A complete set of micro-instructions is defined in section 2.2 and in the next chapter.

## The instructions

Instructions and their execution can be modelled in different ways. The first possibility resembles denotational semantics; a function which interprets the current instruction is constructed. The actions of this function are in close correspondence with the instruction cycle of the CPU. The current instruction is fetched from the memory and the machine state is adjusted accordingly. In order to use pattern matching on the kind of the current instruction at the top level, it must be at a fixed place, therefore the program is usually modelled as a list of instructions. The head of this list is the current instruction. The possibil-
ity we propose here uses a more realistic memory model and the instructions change the state of the machine themselves.

In this description the instructions are functions that take the current machine state as argument and deliver the new state, as changed by the instruction. So an instruction is in the description a function of type state $->$ state. Hence, the general format of an instruction is mnemonic instruction-arguments state $=$ new-state. The state transformation is described in terms of the microinstructions introduced above. Although it is possible to write complex expressions in the instruction specifications, this is not recommended. To obtain a clear specification the new machine state yielded by the instruction is specified using only micro instructions, arguments and constants. To enhance the readability locally defined identifiers are used in the construction of the machine state.

A jump to subroutine instruction with the corresponding return instruction for the machine with a stack as introduced above can be described as:

```
jsr :: address -> instruction
jsr address (pc, reg, st, mem)
    = (pc', reg, st', mem)
    where pc' = address
        st' = s_push pc st
rin :: instruction
ith (pc, reg, st, mem)
    = (pc', reg, st', mem)
        where pc' = s_top st
            st' = s_pop 1 st
```

We prefer to specify the instructions as above instead of the shorter equivalent given below, since it is better readable.

```
jsr :: address -> instruction
jsr address (pc,reg,st,mem) = (address,reg,s_push pc st,mem)
```

A jump to subroutine and corresponding return instruction for the machine with a stack maintained with a stack pointer in the memory are:

```
jsr :: address -> instruction
jsr address (pc, reg, sp, mem)
        = (pc', reg, sp', mem')
        where p\mp@subsup{'}{}{\prime}= address
            sp' = sp_dec sp
            mem'= m_update sp' pc mem
```

```
rtn :: instruction
rtn ( \(p c\), reg, st, mem)
\(=\) ( \(p c^{\prime}\), reg, \(s p^{\prime}\), mem)
    where \(\mathrm{pc}^{\prime}=\) get_num sp mem
    \(s p^{\prime}=s p \_i n c s p\)
```


## Program execution

The instruction sequence to be executed, the program, is also part of the machine. The program is stored in the memory and the instruction to be executed, the current instruction, is addressed by a program counter. It can be executed by extracting the instruction from the memory and applying it to the current state. This clearly reflects the von Neumann machine instruction cycle.

There are two possible places to increment the program counter in order to make the next instruction the current instruction after one instruction cycle. It is possible to increment the program counter in each instruction explicitly, or it can be done in the instruction cycle. Since most machine models use the latter approach we have adopted it here. An instruction cycle executing a single instruction of the machine described above is modelled by:

```
instruction_cycle :: state -> state
instruction_cycle (pc, reg, st, mem)
    = current_instruction (pc', reg, st, mem)
    where pc' = pc_next pc
        current_instruction = get_ins pc mem
```

Note that the function instruction_cycle is not a machine instruction, although it has the same type as an instruction. It is just a function that executes the current instruction that happens to have the same type as instructions.

The instruction cycle above executes just one instruction, in most situations it is more convenient that a machine continuously executes instructions. This can be achieved by applying the instruction cycle recursively to the state delivered by the current instruction.

```
instruction_cycle :: state -> state
instruction_cycle (pc, reg, st, mem)
    = instruction_cycle (current_instruction (pc', reg, st, mem))
    where pc' = pc_next pc
        current_instruction = get_ins pc mem
```

A complete machine description consists of the definition of its state, and the definition of all instructions able to alter this state. The use and purpose of defining the machine are not given by the machine description. Many abstract machines are defined with a very specific purpose, this aim cannot be deduced from a bare machine description, so usually there must be also a description of the objective of the abstract machine to give it any sense.

### 2.2 Description of a simple machine

To compare the functional description with a conventional description we present two descriptions of a simple machine. First, the conventional description will be shown. Then the functional specification is given. In chapter 3 an elaborated abstract machine description is treated.

The machine presented here is a simple conventional machine, called Mac-1. Mac stands for macroarchitecture. It is used by Tanenbaum [Tanenbaum 84] to illustrate the concept of micro-programming.

The Mac-1 is a small machine with a memory of 4096 16-bit words. This memory contains the program to be executed and a stack. The stack grows from high memory addresses to lower ones. The top of the stack is indicated by the stack pointer (sp). The current instruction is indicated by the program counter ( pc ). The machine has one register, the accumulator (ac), to store the result of computations etc. The architecture of Mac-1 can be depicted as:


Fig 2.2 The architecture of Mac-1
Four addressing modes are provided in this machine:
immediate: the operand is specified in the low order bits (8 or 12) of the instruction;
direct: the low-order 12-bits of the instruction are the address of the operand;
indirect: the accumulator contains the address of the operand;
local: the operand is on the stack, the offset is given in the low-order 12-bits of the instruction.
The addressing mode is indicated by the instruction used. There is a very limited set of instructions available:
Load: load the accumulator with the specified operand;
Store: $\quad$ store the contents of the accumulator at the specified address in the memory;
Add: $\quad$ add the operand to the contents of the accumulator, the sum is stored in the accumulator;
subtract the operand from the contents of the accumulator; the result is left in the accumulator;
Jump: $\quad$ set the program counter to the specified address conditional to the contents of the accumulator;
Push: $\quad$ push the specified operand on the stack;
Pop:
Swap:
Sp handling: pop an item from the stack;
exchange the contents of the accumulator and stack pointer; increment or decrement the stack pointer by the 8 -bit constant given in the instruction.
We presume that the reader is sufficiently familiar with conventional machine architectures to understand the description and use of the instructions. The instruction set is specified in the figure below. The meaning of the instructions is given by a Pascal fragment. In these fragments, $m[x]$ refers to memory word $x$.

| Binary | Mnemonic Instruction |  | Meanıng |
| :---: | :---: | :---: | :---: |
| 0000xxmxxxxxxxxx | LODD | Load direct | $a c=m[x]$ |
| 0001 xxxxxxxxxxxx | STOD | Store direct | $m[x]=a c$ |
| 0010xxxxxxxxxxxx | ADDD | Add direct | $a c=a c+m[x]$ |
| $0011 \mathrm{xxKx} \mathrm{\% x} \mathrm{\times x} \mathrm{\times x} \mathrm{\times x}$ | SUBD | Subtract direct | $\mathrm{ac}=\mathrm{ac} \cdot \mathrm{m}[\mathrm{x}]$ |
| 0100xxxxxxxxxxxx | JPOS | Jump positive | If $\mathrm{ac} \geq 0$ then $\mathrm{pc}=\underline{x}$ |
| 0101xxxxxxxxmxyx | JZER | Jump zero | If $\mathrm{ac}=0$ then $\mathrm{pc}=\mathrm{x}$ |
|  | JUMP | Jump | $\mathrm{pc}=\mathrm{x}$ |
| 0111mxxxxxxxxxyx | LOCO | Load constant | $\mathrm{ac}=\mathrm{x}(0 \leq \mathrm{x} \leq 4095)$ |
| 1000xxxxxxxxxxxx | LODL | Load local | $\mathrm{ac}=\mathrm{m}[\mathrm{sp}+\mathrm{x}]$ |
| 1001 mexxxxxxxxxx | STOL | Store local | $\mathrm{m}[\mathrm{x}+\mathrm{sp}]=\mathrm{ac}$ |
| 1010xxxxxxxxxxxx | ADDL | Add local | $a c=a c+m[s p+x]$ |
| 1011 mxxxxxxxxyx | SUBL | Subtract local | $a c=a c-m[s p+x]$ |
| 1100xxxxxxxxxxxx | JNEG | Jump negative | If ac < 0 then $p \mathrm{c}=\mathrm{x}$ |
|  | JNZE | Jump non zero | If $\mathrm{ac} \neq 0$ then $\mathrm{pc}=\mathrm{x}$ |
| 1110xxxxxxxxxxx | CALL | Call procedure | $\mathrm{sp}=\mathrm{sp}-1, \mathrm{~m}[\mathrm{sp}]=\mathrm{pc}, \mathrm{pc}=\mathrm{x}$ |
| 1111000000000000 | PSHI | Push indirect | $s p=s p-1, m[s p]=m[a c]$ |
| 1111001000000000 | POPI | Pop indirect | $m[a c]=m[s p], s p=s p+1$ |
| 1111010000000000 | PUSH | Push onto stack | $\mathrm{sp}=\mathrm{sp}-1, \mathrm{~m}[\mathrm{sp}]=\mathrm{ac}$ |
| 1111011000000000 | POP | Pop from stack | $\mathrm{ac}=\mathrm{m}[\mathrm{sp}] . \mathrm{sp}=\mathrm{sp}+1$ |
| 1111100000000000 | RETN | Return | $p \mathrm{c}=\mathrm{m}[\mathrm{sp}], \mathrm{sp}=\mathrm{sp}+1$ |
| 1111101000000000 | SWAP | Swap ac,sp | $\mathrm{tmp}=\mathrm{ac}, \mathrm{ac}=\mathrm{sp}, \mathrm{sp}=\mathrm{tmp}$ |
| 11111100 yYYYYYYY | INSP | Increment sp | $s p=s p+y(0 \leq y \leq 255)$ |
| 11111110 yyyyyyyy | DESP | Decrement sp | Sp $=$ sp $-\mathrm{y}(0 \leq y \leq 255)$ |

ххххххххххкхх is a 12 -bit machine address; in column 4 it is called $x$.
yyyyyyy is an 8 -bit constant; in column 4 it is called $y$.
The Mac-1 instruction set as descnbed by Tanenbaum.

The functional specification presented below is merely a substitute for the second and last column. The mapping of binary numbers to instructions is not difficult, but omitted for reasons of brevity.

## The state of Mac-1

As shown above the architecture of Mac-1 contains three registers (pc, ac and sp ) and a memory. The state of Mac-1 is fully determined by the contents of the three registers and the memory. In our description the state is given by the tuple (pc,ac,sp,mem).

```
state == (pc,ac,sp,mem)
instruction == state -> state
```


## The micro-instruction level of Mac-1

The registers of Mac-1 contain 16-bit integers. Just like the original description we will not use 16 -bit arithmetic for the description. This is perfect to obtain a small and simple description, but the numbers are a super set of the 16 -bit integers. These registers are modelled by Miranda numbers. When it is required that the content of the registers is limited to 16 -bit integers a tailor made arithmetic must be used. In our description all arithmetical operators must be replaced by a proper micro-instruction. Here, no tailor made micro-instructions are needed for the registers, they can be manipulated just like all other numbers in Miranda.

| pc | $==$ | num |
| :--- | :--- | :--- |
| ac | $==$ | num |
| sp | $==$ | num |

The memory is defined by a small set of micro-instructions. The four access functions defined are:
get a: returns the number stored at address a in the memory;
get_ins a: returns the instructions stored at the indicated address in the memory;
update a n: replaces the contents of memory location with address a with the number $n$;
store prog: stores the complete program prog in the memory.
The memory is formally defined by the abstract type mem.

| abstype | mem |
| :---: | :--- |
| with | get $=$ address $\rightarrow$ mem $\rightarrow$ num |
|  | get_ins $=$ address $\rightarrow$ mem $\rightarrow$ instruction |
|  | update $=$ address $\rightarrow$ num $\rightarrow$ mem $\rightarrow$ mem |
|  | store $=$ program $\rightarrow$ mem |

These micro-instructions are of the same complexity as the Pascal notations, such as $m[x]$, used in the previous description. The informal description and type definitions are sufficient to understand them. Though the definition will be seldomly referenced it is included for the sake of completeness and to show that such an implementation is indeed very simple.

The implementation of the abstract type mem can be very simple. We have chosen the description such that it is an error to use instructions as data and vice versa. This enables the storing of instructions directly in the memory, rather than to use numbers which must be decoded to instructions. the micro-instruction definitions shown here have additional rule altematives in order to give better error messages in erroneous situations. In this way the value of the prototype implementation is increased, these error reporting alternatives are not regarded as a part of the specification which must be obeyed by every implementation.

```
program \(=\) = [word]
address \(==\) num
mem \(==\) [word]
word \(:=\begin{aligned} & \text { I instruction } \\ & \\ & \mathrm{N} \text { num } \\ & \end{aligned}\)
        Undef
get address mem = w_num (get_w address mem)
get_ins address mem = w_ins (get_w address mem)
get_w gaddress \(->\) mem \(->\) word
get_w 0 (w:ws) \(=w\)
get_w \(n\) (w:ws) \(=\) get_w (n-1) ws
get_wn[] = error "illegal memory reference"
w_ins :: word -> instruction
\(w_{\text {_ins }}(1 i)=\mathbf{i}\)
\(\mathbf{w}_{-}\)ins ( \(\mathrm{N} \mathbf{n}\) ) = error "treating a number as instruction"
\(w_{\text {_ ins }}\) Undef \(=\) error "taking a instruction from an undefined word"
w_num :: word -> num
\(w_{\text {_num }}(\mathrm{N} n\) ) \(=\mathrm{n}\)
\(\mathrm{w}_{-}\)num (I i) \(\quad=\) error "treating an instruction as data"
w_num Undef = error "taking a number from an undefined word"
store instructions \(=\) instructions ++ rep (4096-\# instructions) Undef
update address new mem = update_w address ( N new) mem
update_w :: addres \(\rightarrow\) w word \(->\) mem \(->\) mem
update_w 0 new (w:ws) = new.ws
```

update_w n new (w:ws) = w : update_w (n-1) new ws
update_w n new [] = error "illegal memory update"

## The Mac-1 instructions

A direct transformation of the meaning of the instructions given above into the proposed description frame work is:

```
lodd x (pc,ac,sp,me)
= (pc,ac',sp,me)
    where ac' = get x me
stod x (pc,ac,sp,me)
= (pc,ac,sp,me')
    where me' = update }\times\mathrm{ ac me
addd x (pc,ac,sp,me)
= (pc,ac',sp,me)
    where ac' = ac + get x me
subd x (pc,ac,sp,me)
= (pc,ac',sp,me)
    where ac' = ac-getxme
jpos x (pc,ac,sp,me)
= (pc',ac,sp,me)
    where pc' = cond (ac >=0) }\times\textrm{pc
jzer x (pc,ac,sp,me)
= (pc',ac,sp,me)
    where pc' = cond (ac=0) x pc
jump x (pc,ac,sp,me)
= (pc',ac,sp,me)
    where p\mp@subsup{c}{}{\prime}=x
loco x (pc,ac,sp,me)
= (pc,ac',sp,me)
    where ac' = x
lodl x (pc,ac,sp,me)
= (pc,ac',sp,me)
    where ac' = get ( }sp+x)m
stol x (pc,ac,sp,me)
= (pc,ac,sp,me')
    where me' = update (sp + x) ac me
addl x (pc,ac,sp,me)
= (pc,ac',sp,me)
    where ac' = ac + get ( }sp+x)m
subl x (pc,ac,sp,me)
= (pc,ac',sp,me)
    where ac' = ac-get (sp + x) me
jneg x (pc,ac,sp,me)
= (pc',ac,sp,me)
    where pc' = cond (ac<0) xpc
```

```
jnze \(\times(p, a c, s p, m e)\)
```

jnze $\times(p, a c, s p, m e)$
$=\left(p c^{\prime}, a c, s p, m e\right)$
$=\left(p c^{\prime}, a c, s p, m e\right)$
where $\mathrm{pc}^{\prime}=$ cond (ac $\left.\sim 0\right) \times \mathrm{pc}$
where $\mathrm{pc}^{\prime}=$ cond (ac $\left.\sim 0\right) \times \mathrm{pc}$
call $x$ ( $p c, a c, s p, m e$ )
call $x$ ( $p c, a c, s p, m e$ )
$=\left(p c^{\prime}, a c, s p^{\prime}, m e^{\prime}\right)$
$=\left(p c^{\prime}, a c, s p^{\prime}, m e^{\prime}\right)$
where $\mathrm{pc}^{\prime}=\mathrm{x}$
where $\mathrm{pc}^{\prime}=\mathrm{x}$
$\mathrm{sp}^{\prime}=\mathrm{sp}-1$
$\mathrm{sp}^{\prime}=\mathrm{sp}-1$
$m e^{\prime}=$ update $s p^{\prime} p c m e$
$m e^{\prime}=$ update $s p^{\prime} p c m e$
pshi ( $p c, a c, s p, m e$ )
pshi ( $p c, a c, s p, m e$ )
= (pc,ac,sp',me')
= (pc,ac,sp',me')
where $\mathbf{s p}^{\prime}=\mathbf{s p - 1}$
where $\mathbf{s p}^{\prime}=\mathbf{s p - 1}$
$\mathrm{me}^{\prime}=$ update $\mathrm{sp}^{\prime} \mathrm{ac}$ me
$\mathrm{me}^{\prime}=$ update $\mathrm{sp}^{\prime} \mathrm{ac}$ me
pорі (pc,ac,sp,me)
pорі (pc,ac,sp,me)
$=\left(p c, a c, s p^{\prime}\right.$, me' $\left.^{\prime}\right)$
$=\left(p c, a c, s p^{\prime}\right.$, me' $\left.^{\prime}\right)$
where $\mathrm{sp}^{\prime}=\mathrm{sp}+1$
where $\mathrm{sp}^{\prime}=\mathrm{sp}+1$
me' $=$ update ac top me
me' $=$ update ac top me
top $=$ get spme
top $=$ get spme
push (pc,ac,sp,me)
push (pc,ac,sp,me)
$=\left(p c, a c, s p^{\prime}, m e^{\prime}\right)$
$=\left(p c, a c, s p^{\prime}, m e^{\prime}\right)$
where $s p^{\prime}=s p-1$
where $s p^{\prime}=s p-1$
$m e^{\prime}=$ update $s p^{\prime}$ ac me
$m e^{\prime}=$ update $s p^{\prime}$ ac me
pop (pc,ac,sp,me)
pop (pc,ac,sp,me)
$=\left(p c, a c^{\prime}, s p^{\prime}, m e\right)$
$=\left(p c, a c^{\prime}, s p^{\prime}, m e\right)$
where $s p^{\prime}=s p+1$
where $s p^{\prime}=s p+1$
$a c^{\prime}=$ get $s p m e$
$a c^{\prime}=$ get $s p m e$
retn ( $p \mathrm{c}, \mathrm{ac}, \mathrm{sp}, \mathrm{me}$ )
retn ( $p \mathrm{c}, \mathrm{ac}, \mathrm{sp}, \mathrm{me}$ )
$=\left(p c^{\prime}, a c, s p^{\prime}, m e\right)$
$=\left(p c^{\prime}, a c, s p^{\prime}, m e\right)$
where $p c^{\prime}=$ get sp me
where $p c^{\prime}=$ get sp me
$s p^{\prime}=s p+1$
$s p^{\prime}=s p+1$
swap ( $p c, a c, s p, m e$ )
swap ( $p c, a c, s p, m e$ )
= (pc,ac',sp',me)
= (pc,ac',sp',me)
where $a c^{\prime}=s p$
where $a c^{\prime}=s p$
$s p^{\prime}=a c$
$s p^{\prime}=a c$
insp y (pc,ac,sp,me)
insp y (pc,ac,sp,me)
$=\left(p c, a c, s p^{\prime}, m e\right)$
$=\left(p c, a c, s p^{\prime}, m e\right)$
where $s p^{\prime}=\mathbf{s p}+\mathbf{y}$
where $s p^{\prime}=\mathbf{s p}+\mathbf{y}$
desp y (pc,ac,sp,me)
desp y (pc,ac,sp,me)
$=\left(p c, a c, s p^{\prime}, m e\right)$
$=\left(p c, a c, s p^{\prime}, m e\right)$
where $s p^{\prime}=\mathbf{s p - y}$

```
    where \(s p^{\prime}=\mathbf{s p - y}\)
```


## Execution of Mac-1 programs

In order to execute a Mac-1 program it must be stored in the memory. The program execution is chosen to start at address 0 . A real Mac-1 machine must run forever, but a simulation must terminate. The simulation is aborted when the program counter becomes negative. We have chosen to show the program counter, the stack pointer and the contents of the accumulator in the final machine state.

```
run :: program -> (pc,ac,sp)
run program
    = instruction_cycle (pc,ac,sp,mem)
        where pc=0
            ac}=
            sp}=
                mem = store program
instruction_cycle :: state -> (pc,ac,sp)
instruction_cycle (pc,ac,sp,mem)
    = instruction_cycle (instr (pc',ac,sp,mem)) , if pc >=0
    = (pc,ac,sp) ,otherwise
    where pc' = pc+1
            instr = get_ins pc mem
```

It is easy to write an instruction cycle that delivers trace information about the executed program. The trace information shown after each instruction execution could consist of the current contents of the registers and the top of the stack.

Unfortunately it is not possible to show the current instruction in the trace. The instructions stored in the memory of Mac-1 are partially parametrized functions. The Miranda system ${ }^{1}$ shows all partially parametrized functions as <function> and the comparison of such curried functions causes an erroneous program termination with the message program error: attempt to compare functions. This behaviour is suggested by the computational model used; the $\lambda$-calculus. In the $\lambda$-calculus the usual semantics of a comparison of two functions is a test for extensional equality: do these functions have the same reduct for all possible arguments. This is in general an undecidable property.

For curried functions like the instructions of Mac-1 the reduction behaviour of these partially parametrized functions is irrelevant. We are interested in the syntactical value of these expressions. In such a setting two functions are equal if their normal forms are syntactically equal. A test for syntactic equality is part of the semantics of rewrite systems. Such a syntactical comparison can be added to the $\lambda$-calculus, it is known as Church's $\delta$ rule; $\lambda$-calculus with this

[^2]extension is sound, but Church's $\delta$ rule itself is not $\lambda$-definable [Barendregt 84]. Church's $\delta$ rule will not solve our problem. We want to show the function name and not in the corresponding $\lambda$-term which is the normal form of that function name. Moreover, the $\lambda$-term of an recursive function has not a normal form.

The rewrite semantics is obviously the one needed here for the partially parametrized functions. It would be better if the semantics of functional programming languages with pattern matching was based on rewrite systems. (Concurrent) Clean is a functional programming language based on term graph rewriting that allows the suggested use of partially parametrized functions [Brus 87, Eekelen 89].

## An example Mac-1 program

To illustrate the possibility of program execution on the given Mac-1 specification an implementation of the nfib function in Mac-1 code is shown below. The result of this function is equal to the number of function calls needed to compute it. The nfib-number is obtained by dividing the result by the number of seconds needed to compute it, i.e. the number of function calls per second. The Miranda definition of the nfib function reads:

```
nfib 0 = 1
nfib 1 = 1
nfib n = 1 + nfib (n-1) + nfib (n-2)
```

The implemented function nfib expects its argument in the accumulator and also leaves its result in that register. The function implementation is straightforward. The evaluation of the function call nfib 5 on Mac-1 can be simulated by reducing the expression run (nifb 5) in Miranda.

```
nfib :: num -> program
nfib n
= | | (loco 100) , || 0: mnital stack pointer
    I ( swap ) , || 1: load stack ponter
    I (loco n ), || 2: man; load argument for nfib
    | (call 5 ) , || 3: call nfib
    | (jump stop ) , || 4: stop machme
```



```
| (subd 22 ) , || 8. nfib2: compute n-1
    |(fnze 12 ) , || 9 jump to nfib2 if n=1
    | (loco 1 ), || 10: argument = 1; load accumulator with result
    | (retn ), || 11: done, retum to caller
```

|  | ( push |  | ) | \|| 12: nfib3 | save n -1 on the stack |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | (subd | 22 | $)$ | \|| 13: | compute n-2 |
| 1 | ( call | 5 | ) | \|| 14: | compute nfib ( n -2) |
| 1 | ( push |  | ) | \|| 15: | save nfib ( $\mathrm{n}-2$ ) on the stack |
| 1 | ( lodl | 1 | ) | \|| 16: | load $\mathrm{n}-1$ into accumulator |
| 1 | ( call | 5 | $)$ | \|| 17: | compute $\mathrm{nfib}(\mathrm{n}-1$ ) |
| 1 | ( addl | 0 | ) | \|| 18: | add the computed nfib values |
| 1 | ( addd | 22 | $)$ | \|| 19: | add 1 |
| , | ( insp | 2 | ) | \|| 20: | clean up stack |
| 1 | ( retn |  | ) | \|| 21: | done; retum to caller |
|  |  |  |  |  | The constant 1. |

The shown program notation with addresses is rather cumbersome to use. It is more convenient to use an assembler level where labels can be used instead of numbers, also the constructors to distinguish numbers and instructions can be generated by the assembler. In the next chapter such an assembler level is introduced for the ABC-machine described there.

The Miranda interpreter executes circa 20 Mac-1 instructions each second. The resulting nfib-number is 2 . This is not extremely fast, but sufficient to execute small Mac-1 programs. This execution of Mac-1 programs is certainly much faster and more accurate than the pencil and paper simulations necessary with the first specification.

### 2.3 Discussion

The description of the instruction set of Mac-1 in Miranda is a bit longer than the original description given in the table. This is caused by the layout rules we have imposed ourselves to obtain the clearest description. Instructions can also be described in a single line as shown in Section 2.1. Also the binary representation of instructions is not shown in the functional description. Nevertheless, very clear and compact specifications are obtained using this description method. Due to the two level description the operational semantics of each instruction are clearly specified in terms of simple micro-instructions.

The micro-instructions are simple access functions defining the machine components. By adjusting the micro-instruction level any amount of detail can be incorporated in these machine components. The machine description shown here can be extended to perform 16-bit arithmetic and allow the use of instructions in the memory as data and vice versa by changing only the microinstruction level.

The use of a functional programming language as description formalism has a number of advantages:

- The description can be verified by the compiler for the functional language and has well defined semantics.
- The specification is more complete. For instance, in the functional description it is clear when the program counter is incremented, and that the incremented program counter is stored on the stack in a call instruction.
- The description is more direct. There are no temporary variables needed, as used in the imperative description of the swap instruction. This is an advantage of having an explicit state in the description of the instructions instead of an implicit global state.
- The description serves as a prototype implementation of the machine and is able to exccute programs. When programs execute correctly on the prototype, this increases the confidence in the correctness of the description and the suitability of the specified machine.
It is a pity that Miranda lacks rewrite semantics, this limits the use of the specification to trace the execution of programs.

In order to use the specification for the execution of large programs an assembler level is needed.

The functional descriptions require, just like many other description methods, a single locus of control and a tree-like structure for the state. When a machine does not possess this structure the actual structure must be described. For instance, a parallel machine description contains a scheduler to model the concurrent execution. In this way a parallel architecture is modelled by a sequential machine and the description method can be applied, although there is a lot more description overhead.

## Chapter 3 The ABC-Machine

3.1 The architecture of the ABC-machine
3.2 The micro-instruction set
3.3 The instruction set
3.4 Program execution
3.5 ABC-assembler
3.6 Comparison with the G-machine
3.7 Discussion

This chapter presents a successful application of the operational machine description method introduced in the previous chapter for a large and complex concrete problem. The development of this specifcation was a test of the suitabilty of the description method. In order to be useful this description method must yield specifications which are: a clear definition for the product to construct; a useful prototype implementation; fast and easy to develop.

The ABC-machine is an imperative abstract machine architecture for graph rewriting. To enable a detailed description of the graph rewriting process a huge number of instructions is defined. Circa 50 instructions are used to describe the rewriting process, even a larger number of instructions is provided to handle basic values. The letters A, B and C represent the three stacks used in this machine. This virtual machine is designed to describe graph reduction on a low level of abstraction, close to the level of a concrete machine architecture. This machine is an abstraction of the large class of concrete stack based machine architectures. ABC-machine code is used as intermediate level on the compilation path of functional languages to concrete machine code.

The first step on this translation path consists of the translation of functional languages to the graph rewriting language Clean [Eekelen 89]. In this translation step the syntactical sugar of functional languages is removed and sharing of computations is specified accurately. In many cases the translations from a functional program to an equivalent Clean program are straightforward.

In some cases, like guards and ZF-expressions, the transformations become tedius and rather complex. These transformations have also been formally specified in a functional language [Koopman 88]. In the next step Clean is translated to ABC-code. This translation links the world of graphs and strategies to a sequence of imperative instructions and is described in chapter 5. Finally, ABCinstructions are translated to, or interpreted by, concrete machine code. During this last step all details of the concrete machine become important to achieve an efficient implementation. A simple implementation of the ABC-machine is relatively easy to make, for instance using macro expansion.


Fig 3.1 The place of the ABC-machine in the translation process of functional languages.
There are several reasons to introduce the additional intermediate level defined by the abstract machine on the compilation path. First of all, the abstract machine helps to get a more structured implementation. When the abstraction is well designed, irrelevant machine dependent issues are omitted. The trade-offs in the imperative graph rewriting process are shown clearly. Some of the restrictions imposed by a concrete machine, such as limited resources (like the limited stack size) are not necessarily present in an abstract machine. Of course, all problems have to be solved on one level or the other. However, the separation of concerns helps to get a more structured view on the problems (and their solutions) one has to face when making an implementation of functional languages.

Another advantage of the additional intermediate level is that the portability of the implementation is increased. To implement functional languages on a new concrete machine only a new ABC-machine has to be implemented. Since the ABC -machine is an imperative machine on a relatively low level of abstraction, it is much easier to implement directly than Clean or a high level functional language.

It is not feasible to use an existing imperative programming language (like C or Pascal) instead of ABC-code since it imposes too much restrictions on the language constructs possible, hence the execution speed will be too low. So, a tailor made abstract machine is developed.

In the next section a global overview of the architecture of the ABC-machine is given by specifying its state. In section 3.2, the basic compo-
nents of the machine are introduced by the definition of the micro-instruction set. The actual ABC-machine instructions are specified in terms of these microinstructions in section 3.3. How the abstract machine executes an ABC-program is explained in section 3.4. In section 3.5 an assembler form of the ABC -instructions is introduced to improve the readability and to simplify the generation of ABC-programs. In section 3.6 the ABC-machine is compared with the well-known G-machine [Johnson 84], the machines themselves and the description methods are both contrasted. Finally, the ABC-machine and its description method are briefly discussed.

### 3.1 The architecture of the ABC-machine

The main part of the instruction set of the ABC-machine is designed to conveniently express graph rewriting. When this instruction set is mapped on a traditional concrete machine the resulting code will be executed relatively slow, since present-day architectures are not at all designed for graph rewriting. A graph structure is not directly available on these architectures, but it has to be represented by some data structure. The modification of the graph will lead to complex memory management problems, involving garbage collection. For many simple calculations the use of graphs is not effective. This observation has triggered the introduction of the second part of the abstract machine. This part is an abstraction from a traditional stack-based architecture. To obtain an efficient program these fast instructions are used and graph rewriting is avoided whenever possible!

The ABC-machine consists of the following memories:

- the $A($ rgument $)$-stack used to reference nodes in the graph store;
- the B(asic value)-stack used to hold and manipulate basic values efficiently;
- the C(ontrol)-stack to store and retrieve return addresses;
- the graph store containing the graph to be rewritten;
- the descriptor store containing information about the symbols used;
- the program counter contains an identification of the instruction (instrid) to be executed;
- the program store containing the instruction sequence to be executed;
- an $i(n p u t)$-o(utput) channel to enable interaction with the world.

This machine can be depicted ${ }^{1}$ as:

[^3]

Fig 3.2 The ABC-machine
The corresponding state of the ABC-machine is defined by the following tuple:

$$
\begin{aligned}
\text { state }== & \text { (astack, bstack, cstack, graphstore, descstore, instrid, programstore, io) } \\
& \text { II This state is denoted in instructions as: (as,bs,cs,gs,ds,pc,ps,10) }
\end{aligned}
$$

The next section contains a thorough specification of these components.

### 3.2 The mlcro-instruction set

The components of the ABC-machine are abstract data-types. They are sufficiently specified by the operations defined on them. The Miranda implementation of the micro-instructions is not presented here, but it can be found in appendix A. These Miranda definitions specify the semantics of the microinstructions, on the other hand they suggest too much a specific implementation. The given implementation must be viewed as a definition of the semantics, not as a suggestion for an efficient implementation. It is merely needed for detailed information like the index of the top of a stack.

The following type synonyms are used to increase the clarity of the type definitions. The type nat stands for natural numbers including zero, it is represented by Miranda numbers (num).

| arity | nat |  |
| :---: | :---: | :---: |
| a_dst | nat | Il the index of the destination on the A-stack |
| a_src | nat | II the index of the source on the A-stack |
| b_dst | nat | Il the index of the destination on the B-stack |
| b_src | nat | II the index of the source on the B-stack |
| c_src | nat | II the index of the source on the C-stack |
| nr_args | nat | II the number of arguments involved |
| arg_nr | nat | II the number of the argument involved |

To enhance the value of the specification in appendix A as prototype implementation, a number of additional rule alternatives are used to create sensible error messages.

## The graph store

The graph in the graph store is basically a Clean-like graph. It is composed of nodes. Each node is labelled with an unique identification, a node-id.

There are micro-instructions to generate a new (empty) graph store, to create a new (empty) node in the graph and to get a node out of the graph. Finally, a node can be updated by a function passed as parameter to the gs_update micro-instruction.

```
gs_get = nodeid -> graphstore -> node
gs_וnit = graphstore
gs_newnode = graphstore -> (graphstore,nodeld)
gs_update : nodeid -> (node >> node) -> graphstore -> graphstore
```

Note that each of these micro-instruction names starts with gs_ (for graph store).

## The nodes

Each ABC-node contains a desc-id (descriptor identification) to represent the symbol in the Clean-node. A desc-id is an entry in the descriptor store indicating the symbol, it is represented by the type descid. Furthermore, a node contains a sequence of node-id's representing the arguments. A node can also contain a basic value like an integer or boolean. To reduce the number of definitions only these two basic types are considered, so characters, reals and strings are not treated here.

The graph in the graph store will be examined by the program which will try to reduce the graph to normal form. For this purpose it is convenient that additional information is stored in the nodes of the graph. A node contains a context; the instr-id (instruction identification) referring to the first instruction of an instruction sequence. By convention, this ABC-instruction sequence will reduce the corresponding node to root normal form when it is executed. The instr-id stored in a node can be changed during reduction. This is used for several markings of the node. For instance, to determine at run-time that a node is already under reduction, in this way cyclic computations can be detected.

There are micro-instructions to extract the pieces of information stored in a node. The nodes can be extracted of the graph store by the micro-instruction gs_get defined above.

```
n_arg \(\quad=\) node \(->\) arg_nr \(->\) arity \(\rightarrow\) nodeid
n_args \(\quad=\) node \(->\) arity \(\rightarrow\) nodeid_seq
n_arity \(\quad=\) node \(\rightarrow\) arity
n_B \(\quad=\) node \(->\) boolean
n_descid \(=\) node \(\rightarrow\) descid
```

| n_entry | $:: ~ n o d e ~$ |
| :--- | :--- |
| n_I | instrid |
| n_I | $\because$ node $->$ int |
| n_nargs | $=$ node $->$ nr_args $->$ arily $->$ nodeid_seq |

There are also micro-instructions to test whether a piece of a node has a certain value. These micro-instructions are important for pattern matching.

| n_eq_arity | $=$ node $\rightarrow$ arity $\rightarrow$ boolean |
| :--- | :--- |
| n_eq_B | $=$ node $\rightarrow$ boolean $\rightarrow$ boolean |
| n_eq_descid | $=$ node $->$ descid $\rightarrow$ boolean |
| n_eq_I | node $\gg$ int $\rightarrow>$ boolean |
| n_eq_symbol $::$ node $\rightarrow$ node $>$ boolean |  |

Finally, the contents of a node can be changed by:

| n_copy | $=$ node $\rightarrow$ node $\rightarrow$ node |
| :--- | :--- |
| n_fill | $=$ descid $->$ instrid $\rightarrow$ args $\rightarrow$ node $\rightarrow$ node |
| n_fill | $=$ descid $\rightarrow$ instrid $\rightarrow$ boolean $\rightarrow$ node $\rightarrow$ node |
| n_fill | $\because$ descid $\rightarrow$ instrid $\rightarrow$ int $\rightarrow$ node $\rightarrow$ node |
| n_setentry | $=$ instrid $\rightarrow$ node $\rightarrow$ node |

The micro-instructions to change a node are passed as argument to gs_update to overwrite nodes in the graph.

The micro-instructions above show that the graph store deals with variable sized nodes. Overwriting a previously created node with new values is always possible. In general, one first has to create a new empty node which has to be filled later. This method is in particular convenient for the creation of cyclic graphs. The code generation schemes presented in the chapter 5 rely on the presence of a node that can be overwritten with the result of a reduction.

## The descriptor store

The ABC-machine contains a piece of memory where symbol descriptors are stored. This descriptor store contains information about the symbols used in the rewrite system. Given the descriptor identification, descid, a descriptor can be taken from the descriptor store. This store can be initiated by passing a list of descriptors to the ds_init micro-instruction.

```
ds_get = descid -> descstore -> desc
ds_init = [desc] -> descstore
```


## The descriptors

A descriptor contains information of the associated symbol; its arity, the start address of the reduction code and the name of the symbol. The symbol name is only used to print a representation of the graph in root normal form on the
output channel. Information can be retrieved from the descriptor by the corresponding micro-instructions.

| d_ap_entry | $=$ desc $->$ instrid |
| :--- | :--- |
| d_arity | $=$ desc $\rightarrow$ arity |
| d_name | $=$ desc $->$ string |

## The A-stack

The $A$-stack contains node-id's; references to nodes stored in the graph store. It is used to access the actual arguments and the result of the rewrite rule executed. Just as in other imperative languages, not a complex data structure is passed to or returned from a function, but a reference to these objects is passed. The top of the stack has index 0 (as is the case for the other stacks).

A new, empty, A-stack can be created by:

$$
\text { as_init } \quad=\text { astack }
$$

An element at any depth, or a sequence of nr_args top elements, can be retrieved from the A-stack by:

| as_get | $:$ a_src -> astack -> nodeid |
| :--- | :--- |
| as_topn | $=$ nr_args $->$ astack $->$ nodeid_seq |

The type nodeid_seq represents a sequence of node-id's which can be taken from or pushed on the A-stack, it serves also as the argument sequence for a node. The A-stack can be updated by:

$$
\begin{array}{ll}
\text { as_popn } & =\text { nr_args }->\text { astack }->\text { astack } \\
\text { as_push } & =\text { nodeid }->\text { astack }->\text { astack } \\
\text { as_pushn } & =\text { nodeid_seq }->\text { astack }->\text { astack } \\
\text { as_update } & =\text { a_dst }->\text { nodeid }->\text { astack }->\text { astack }
\end{array}
$$

## The B-stack

If the calculation of a simple numerical value would be performed by building nodes, performing redirections and the like, it is obvious that in no way an efficient implementation can be achieved on traditional hardware which has no support for these kind of actions whatsoever. On a traditional machine simple calculations are performed on a stack using only a couple of simple instructions. In order to obtain the desired behaviour for these calculations, the ABC-machine is equipped with a stack to hold basic values.

A further optimization is the use of registers to hold these basic values. However, registers are not included in the ABC-machine design. The number of registers and the operations possible on them varies too much between different
concrete machines. Register allocation is left as a part of the implementation of the ABC -machine.

The $B$-stack contains basic values like integers and booleans. Such values are stored untagged on the B-stack. This means that it is impossible to determine the type of the elements of the B-stack.

The initial, empty, B-stack is created by:

$$
\text { bs_init } \quad=\text { bstack }
$$

Information can be obtained from the B-stack with the micro-instructions:

```
bs_get = b_src -> bstack -> basic
bs_getB := b_src -> bstack -> boolean
bs_getI = b_src -> bstack >> int
```

The B-stack can be changed with the micro-instructions:

| bs_popn | $:$ b_src -> bstack -> bstack |
| :--- | :--- |
| bs_push | $::$ basic $->$ bstack $->$ bstack |
| bs_pushB | $=$ boolean -> bstack $->$ bstack |
| bs_pushi | $::$ int $->$ bstack $->$ bstack |
| bs_pushn | $=$ basic_seq $->$ bstack $->$ bstack |
| bs_update | $=$ b_dst -> basic -> bstack $->$ bstack |

Beside these updating micro-instructions there are micro-instructions defined to perform computations with the basic values stored on the B-stack. Usually the arguments are all on top of the $B$-stack and are replaced by the result of the operation. When arguments are not on the B-stack they are an argument of the micro-instruction. Some micro-instructions to handle integer values are:

| bs_addI | $=$ bstack $->$ bstack |
| :--- | :--- |
| bs_eqI | $=$ bstack $->$ bstack |
| bs_eqIi | int $->$ b_src $->$ bstack $\rightarrow$ bstack |
| bs_gII | $=$ bstack $\rightarrow$ bstack |

## The program

The ABC-machine contains a sequence of machine instructions representing the reduction algorithm: the program. This program will rewrite the initial graph to its normal form according to the annotated functional strategy. Conceptually there are two algorithms involved, the annotated functional reduction strategy which indicates the next redex and the rewriting of that redex according the Clean rules. These algorithms are merged in the ABC-program to increase the efficiency. Each Clean rule is translated into a sequence of ABC-instructions. This instruction sequence controls the order of reductions required for this rule and performs the reduction according to the first matching rule alternative.

Each instruction has an unique identification, the instr-id, in order to be indicated as the current instruction by the program counter. A program does not change during execution, it is loaded in the machine when the machine is booted. Programs to be stored are denoted by a list of instructions.
ps_get = instrid $->$ programstore $->$ instruction
ps_init $=$ [instruction] $\rightarrow$ programstore

## The program counter

Since the ABC-machine has an imperative nature it is essential to have a locus of control; a program counter. The program counter contains the instr-id of the next instruction to be executed.

There are micro-instructions provided to initiate the program counter (points to the first instruction of the program) and to change it. The program counter can be incremented (i.e. the next instruction becomes the current instruction), or can be set outside the program to indicate that the program is finished. Finally, one can check whether the last instruction of the program is reached.

| pc_init | $=$ instrid |
| :--- | :--- |
| pc_next | $=$ instrid $->$ instrid |
| pc_halt | $=$ instrid $->$ instrid |
| pc_end | $=$ instrid $\rightarrow$ bool |

## The C-stack

The C-stack (Control stack) is used to implement nested reductions in the abstract machine. The program counter can be stored and recovered from this stack.

| cs_init | $=$ cstack |
| :--- | :--- |
| cs_get | $=$ c_src $\rightarrow$ cstack $->$ instrid |
| cs_popn | $=$ nr_args $\rightarrow$ cstack $->$ cstack |
| cs_push | $=$ instrid $\rightarrow$ cstack $\rightarrow$ cstack |

## The input-output channel

The abstract machine furthermore contains an input output channel used to show the result of the reduction to the world outside. On the output channel strings can be printed. These strings are appended to the existing output channel.

```
io_init = io
io_print = string \(\rightarrow\) io \(\rightarrow\) io
```


### 3.3 The Instruction set

Each machine instruction of the ABC-machine is defined in terms of the microinstructions. Each instruction consists of an instruction identifier and zero or more operandi. Not all instructions are shown in this section. The specification in terms of micro-instructions is only given for the most important instructions. For some other instructions only the type and an informal explanation is given. A precise definition of all machine instructions is given in appendix A .

The instructions shown here are classified according to their main purpose:

- graph manipulation;
- retrieving information from a node;
- manipulating the A-stack;
- manipulating the B-stack;
- changing the flow of control;
- generating output.


## Graph manipulation

There are several kinds of instructions to manipulate the graph store. There is only one instruction to create a new node in the graph store. Several instructions can be used to change the contents of existing nodes. Finally, there are also instructions that fetch information stored in the nodes of the graph. All instructions for graph manipulation (with exception of the instruction create) have as operand an offset in the A-stack, to find the node-id of the node to manipulate.

The instruction create creates a new empty node in the graph store, the node-id of the new node is pushed on the A-stack. It is defined as:

```
create (as,bs,cs,gs,ds,pc,ps,io)
\(=(a s ', b s, c s, g s, d s, p c, p s, i o)\)
    where as' = as_push nodeid as
        (gs', nodeid) = gs_newnode gs
```

create is the only instruction to create a new (empty) node. All other instructions can only change the contents of an already existing node. The most elaborated instruction to update the contents of a node is fill, which is defined as:

```
fill desc nr_args instrid a_dst (as,bs,cs,gs,ds,pc,ps,io)
\(=\) (as',bs,cs,gs',ds,pc,ps,io)
    where as' = as_popn nr_args as
        gs' \(\quad=\mathrm{gs}\) update nodeid ( \(n\) _fill desc instrid args) gs
        nodeid \(=\) as_get a_dst as
        args = as_lopn nr_args as
```

As can be seen from this specification the arguments of the node are taken from the A-stack. The first argument is on top of the A-stack.

Other instructions to change the contents of an existing node are:

| fill_a | = a_src -> a_dst -> instruction | Il the copy node instruction; |
| :---: | :---: | :---: |
| fill | $=$ bool $\rightarrow$ a_dst $->$ instruction | Il fills the node with the given boolean; |
| fill ${ }^{\text {dilb }}$ | :- b_src -> a_dst -> instruction | Il fills with a boolean found at the B-stack; |
| filli | : int $->$ a_dst $->$ instruction | Il fills the node with the given integer; |
| fill_b | z b_src -> a_dst -> instruction | II fills with the integer found at the B-stack |
| set_entry | $=$ instrid $->$ a_dst $->$ instruction | II changes the reduction context of the node. |

Building a Clean-node in the graph store involves at least two instructions: a create, which leaves a node-id on the A-stack and one of the fill instructions.

Example: to create the graph Cons 1 Nil, the following instruction sequence can be used. Assume that "_rnf" indicates some ABC-instruction sequence, probably just containing a rtn instruction since the newly created nodes are already in root normal form. The descriptorid's of Nil and Cons are indicated by "Nil" and "Cons". This code fragment is written in the ABC-assembly language introduced below.

| [ | Create |  | 11 node for Cons |
| :---: | :---: | :---: | :---: |
|  | Create |  | Il node for Nil; $2^{\text {nd }}$ arg of Cons |
|  | Fill | "Nil" 0 "_rnf" 0 | Il fill node just created |
|  | Create |  | Il node for $1 ; 1^{\text {St }}$ arg of Cons |
|  | Fill | 10 | 11 fill node just created |
|  | Fill | "Cons" 2 "_mp 2 | II fill Cons node |

## Retrieving information from a node

The retrieved information is stored on one of the stacks or in the program counter. The node-id of the node is found at the indicated depth on the A-stack. There are also instructions to test whether the contents of a node has the given value.

```
push_args a_src arity nr_args (as,bs,cs,gs,ds,pc,ps,io)
\(=\left(a s^{\prime}, b s, c s, g s, d s, p c, p s, i o\right)\)
    where as' = as_pushn args as
    args = n_nargs (gs_get nodeid gs) nr_args arity
    nodeid \(=\) as_get a_src as
pushl_a a_src (as,bs,cs,gs,ds,pc,ps,io)
\(=(a s, b s ', c s, g s, d s, p c, p s, i o)\)
    where bs' = bs_pushl int bs
    int \(=n \_1\) (gs_get nodeid \(g s\) )
    nodeid = as_get a_src as
```

```
eqI_a int a_src (as,bs,cs,gs,ds,pc,ps,io)
\(=(a s, b s, c s, g s, d s, p c, p s, i o)\)
    where bs' = bs_pushB equal bs
        equal \(=\) n_eq_I (gs_get nodeid gs) int
        nodeid = as_get a_src as
eq_desc_arity descid arity a_src (as,bs,cs,gs,ds,pc,ps,io)
\(=\) (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_pushB equal bs
        equal = (n_eq_descid node descid) \& (n_eq_arity node arity)
        node \(=\) gs_get nodeid gs
        nodeid = as_get a_src as
```

In appendix A some other instructions to extract (a selection of) the arguments from a node, change the arguments or to test on other basic value types are given.

## Manipulating the A-stack

The A-stack is used to access the nodes involved in a rewriting. Via push instructions new node-id's can be pushed on the stack. In addition the following instructions are provided to manipulate the A-stack:

```
pop_a nr_args (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_popn nr_args as
push_a a_src (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_push nodeid as
                        nodeid = as_get a_src as
update_a a_src a_dst (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_update a_dst nodeid as
    nodeid = as_get a_src as
```

Example: the cyclic graph ones: Cons 1 ones, is constructed similar to the previous example. Here the advantages of a separate create and fill instruction are visible.

| [ | Create |  | 11 node for Cons |
| :---: | :---: | :---: | :---: |
|  | Push_a | 0 | - $112^{\text {nd }}$ arg of Cons |
|  | Create |  | , Il node for $1 ; 1^{\text {st }}$ arg of Cons |
|  | Filli | 10 | , \|l fill node just created |
|  | Fill | "Cons" 2 "_rnf 2 | ] II fill Cons node |

## Manipulating the B-stack

For the B-stack the same stack handling instructions as for the A-stack are defined:

| pop_b | $:=$ | nr_args $->$ instruction |
| :--- | :--- | :--- |
| push_b | $=$ | $b \_$src $\rightarrow$ instruction |
| update_b | $=$ | b_src $->$ b_dst $->$ instruction |
|  |  |  |
| pushi | $=$ int $\rightarrow$ instruction |  |
| pushB | $=$ bool $->$ instruction |  |

There are also instructions to manipulate the basic values on this stack. A typical example is

```
addI (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_addl bs
```

There are many more instructions to do arithmetic, they all follow the same scheme as the add instruction presented here. They are not listed here for reasons of brevity.

## Changing the flow of control

The desired flow of control has to be realized by manipulating the program counter. Jumps are unconditional, or directed by the boolean value on top of the B-stack.

```
jmp address (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs,cs,gs,ds,pc',ps,10)
    where \(\mathrm{pc}^{\prime}=\) address
jmp_false address (as,bs,cs,gs,ds,pc,ps,io)
\(=\) (as,bs',cs,gs,ds,pc',ps,io)
    where pc' \(=\) cond (bs_getB 0 bs) pc address
        bs' = bs_popn 1 bs
jmp_true address (as,bs,cs,gs,ds,pc,ps,io)
    \(=\left(a s, b s s^{\prime}, c s, g s, d s, p c, p s, i o\right)\)
    where pc ' \(=\) cond ( bs _getB 0 bs ) address pc
        bs' = bs_popn 1 bs
```

When a jsr (jump to subroutine) instruction is executed, the current value of the program counter is stored on the C -stack, a rtn (return from subroutine)
instruction will restore the program counter and pop the return address from the C -stack.

```
Jsr address (as,bs,cs,gs,ds,pc,ps,10)
\(=\) (as,bs,cs',gs,ds,pc',ps,io)
    where \(p c^{\prime}=\) address
        Cs' = cs_push pccs
rtn (as,bs,cs,gs,ds,pc,ps,10)
= (as,bs,cs',gs,ds,pc',ps,ıo)
    where pc ' \(=\) cs_get 0 cs
        cs' \(=\) cs_popn 1 cs
```

A jsr_eval instruction will start the execution of the code addressed in the node referenced by the top of the A-stack, the return address is saved on the C-stack. Hence, it performs a jsr to the address stored in the node. By convention, executing the instruction sequence will reduce the node to its root normal form.

```
jsr_eval (as,bs,cs,gs,ds,pc,ps,ıo)
= (as,bs,cs',gs,ds,pc',ps,o)
    where pc' = n_entry (gs_get nodeld gs)
        nodeid = as_get 0 as
        cs' = cs_push pccs
```

The program execution stops after the execution of a halt instruction. The instruction fetch_cycle (see below) will be left when the program counter contains an address outside the program.

```
halt (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs,gs,ds,pc',ps,ı0)
    where pc' = pc_halt pc
```


## Generating output

To show the result of the reduction there are print instructions. These instructions append strings to the output channel.

```
prınt string (as,bs,cs,gs,ds,pc,ps,ıo)
= (as,bs,cs,gs,ds,pc,ps,10')
    where 10' = 10_print string to
print_symbol a_src (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs,gs,ds,pc,ps,10')
    where io' = io_print string io
        node = gs_get (as_get a_src as) gs
        string = symbol_to_string node desc
        desc = ds_get (n_descid node) d
```


### 3.4 Program execution

To run the ABC-machine described here, the instructions must be applied to the state of the machine. Once started the machine must continue to execute instructions until the halt instruction is executed. While the machine is running, the current instruction is fetched continuously out the program store and applied to the present state.

## The instruction cycle

The fetch cycle recursively fetches the current instruction out the program and applies it on the current state. The fetching (and hence the machine) stops when the program counter indicates that a halt instruction is executed.

```
fetch_cycle :: state -> state
fetch_cycle (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs,gs,ds,pc,ps,io), pc_end pc
= fetch_cycle (currinstr (as,bs,cs,gs,ds,pc',ps,io)), otherwise
    where pc' = pc_next pc
    currinstr = ps_get pc ps
```


## Booting the machine

Before it is possible to run the machine, the machine is loaded (booted) with the initial state. The program and descriptors must be supplied as argument to the boot function. All parts of the machine are initiated by the corresponding init micro-instructions. The machine starts evaluating the program on the first instruction.

```
boot :: ([instruction],[desc]) -> state
boot (program,descriptors)
= (as,bs,cs,gs,ds,pc,ps,io)
    where pc = pc_init
        as = as_init
        bs = bs_init
        CS = cs_init
        gs = gs_init
        ps = ps_init program
        io = io_init
        ds = ds_init descriptors
```


### 3.5 ABC-assembler

In order to write ABC-programs there must be a denotation for all objects involved. The denotation of the function arguments is specified here.

| boolean | $==$ | bool II A boolean is represented by the Miranda type bool. |
| :--- | :--- | :--- |
| context | $==$ | instrid |
| II A context is represented by the address of the first instruction. |  |  |
| descid | $==$ | instrid II A descriptor-id is the address of the descriptor in the store. |
| instrid | $==$ | nat |
| int An address is the number of the instruction in the sequence. |  |  |
| nat | $==$ | num II Integers are represented by Miranda num's |
|  | $==$ int | II Natural numbers are also represented by Miranda num's. |

References to instructions are made by the address of the instruction within the program. This is fine for the ABC-machine and close to reality in a concrete machine, but cumbersome to read and write for human beings. Moreover the standard Miranda system does not support functions on top level very well. The system just prints <function>, which is not a clear identification of a specific ABC -instruction. To make ABC-programs better understandable an assembler level is introduced. Assembler statements can be mapped directly to instructions, but the assembler uses labels instead of addresses. A program in ABC-assembler is represented by a Miranda data structure.

| assembler <br> label <br> desc_label <br> red_label |  | [statement] [char] label label | II the label of a descriptor II the label of an instruction |
| :---: | :---: | :---: | :---: |
| statement |  |  |  |
|  | := | Label | label |
|  |  | Descriptor | desc_label red_label arity name |
|  |  | Create |  |
|  |  | Fill | label nr_args label a_dst |
|  |  | Jmp | label |
|  |  | etc... |  |

Every ABC-instruction is represented by a constructor with a similar name, the first character is changed in an uppercase character to make it a constructor. Generally also the arguments are identical, only addresses are replaced by labels. ABC-assembler also contains labels and descriptor definitions. See appendix A for a complete definition of ABC -assembler statements and an assembler.

Example: this assembler form to represent ABC-instructions allows us to write:

```
example :: assembler
example
\(=\left[\begin{array}{ll}\text { "Length1" } & \text { Jsr_eval }\end{array}\right.\)
            Eq_desc_arity "Cons" 00 ,
            Jmp_false "Length2" ,
            Push_args 02
            Create
                Push_a 2
                Fill
                                "Length" 1 "n_Length" 1
```

Instead of the corresponding ABC-instruction sequence:

```
example :: program
```

example

| = [ | jsr_eval |  |
| :---: | :---: | :---: |
|  | eq_desc_arity | 2300 |
|  | jmp_false | 62 |
|  | push_args | 02 |
|  | create |  |
|  | push_a | 2 |
|  | fill | 1216131 |

The reason for introducing a separate assembler level is not only to enhance the readability of the program, but also to simplify the translation of Clean rules to ABC-code. This simplification is caused by the use of symbolic names and labels instead of the corresponding numbers of the ABC-machine. To have ABC-assembler as a Miranda data-structure has also some advantages for the prototype; it eliminates the generation and parsing of ABC-assembler represented as a list of characters and enables the printing and manipulating of statements which is impossible for a sequence of instructions.

### 3.6 Comparison with the G-machine

Another well known abstract machine for graph reduction is the G-machine [Johnson 84]. The objectives of this machine are identical to that of the ABC-machine: defining an imperative abstract graph rewrite machine as an intermediate level in the compilation of functional languages. The G-machine is not described in a functional language, and is initially meant for graphs in the applicative form (each argument is bound to the function by an AP node). Due to these two reasons both the description and the described abstract machines differ in almost all details. An overall view of the ABC-machine and the G-machine shows many correspondences.

In order to compare the G-machine and its description with the ABC-machine, the G-machine state and some instructions in the G-machine formalism are shown. Afterwards the main differences between description methods and abstract machines are indicated. The description method used for the G-machine very much resembles the description method used by Landin to describe the SECD-machine [Landin 64].

## The G-machine state

A state in the abstract G -machine is a 7 -tuple $\langle O, C, S, V, G, E, D>$, where
O is the output produced so far.
C is the G-code sequence currently being executed (sometimes called codestack).
S is a stack of node names i.e. pointers into the graph.
V is a stack of basic values.
G is the graph. A mapping from node-names to nodes. There are nodes of the following type:
INTi integer nodes,
BOOLb boolean nodes,
NIL empty-list nodes,
CONS $n_{1} n_{2}$ list nodes, where $n_{1}$ is the pointer to the head graph and $n_{2}$ is the pointer to the tail graph,
AP $n_{1} n_{2}$ application nodes, where $n_{1}$ is a pointer to the function graph and $n_{2}$ is a pointer to the argument graph,
FUN $f \quad$ a node with a reference to the compiled function $f$,
HOLE a node which is to be filled with another value later; it is used during the construction of cyclic graphs.
E is the global environment, which is a mapping from function names to pairs consisting of the number of curried arguments of the function and its code sequence.
D is a dump used for recursive calls to EVAL: a stack of pairs consisting of; a stack of node names: S before EVAL, and a G-code sequence: C before EVAL.

## Comparison of machine states.

The correspondence of the machine states is shown by indicating which components of the ABC-machine correspond to the elements of the 7 -tuple of the G-machine described above.
O The output has a direct correspondence with the output in the ABC-machine.
C In the G-machine the current instruction is the head of the code sequence C . The flow of control involves the movement of code between the code sequence C , the environment E , and the dump D . In the ABC -machine there
is no movement of code; the instruction sequence executed is controlled by the program counter (pc) and the control stack (C-stack).
S The stack, S, of the G-machine directly corresponds to the argument stack (A-stack) of the ABC-machine.
$\vee$ The stack of basic values, V , corresponds to the basic value stack ( B -stack), although it is used in a more restricted way in the G-machine.
G The graphs of both machines obviously correspond. The G-machine contains only binairy nodes while the nodes in the ABC-machine are variable sized. Moreover, the nodes of the ABC -machine contain an extra field for the context and information about the constructor can be found via the descriptor-id identifying the symbol of the node. The type list is the only data structure in the G-machine, it is "wired" in the definition. In the ABC-machine there are no assumptions on the data types used, the corresponding constructors are sufficiently specified by their descriptors in the descriptor store.
E The environment, E , of the G-machine contains the arity and the code corresponding to the functions. The arity is stored in the descriptor store of the ABC-machine, and the code in the program.
D The dump of the G-machine is not found in the ABC-machine. The stack is not stored separately on an other stack in the ABC-machine, the actions needed (number of pops upon returning from a function) are computed during the compilation of Clean to ABC-code. The code sequences stored in the dump of the G-machine, are recorded by the program counters stored on the C-stack of the ABC-machine.

## The G-machine instructions

In a G-machine state, () denotes an empty stack or an empty code sequence. The initial graph is indicated by 13 . The semi-colon appends values onto an output sequence. The period is used as infix cons for instruction sequences and push for stacks. Updating of the graph is written as $\mathrm{G}\left[\mathrm{n}=\mathbb{N} \mathrm{N}_{1}\right]$. If there is a node named $n$ previously in G, then the node $n$ is updated with a new value, otherwise a new node is created. This notation is also used in pattern matching situations.

The flow of control is conducted by the following instructions. Descriptions are taken from [Johnson 84], with a few errors corrected (see also [Peyton Jones 87]).

$$
\begin{aligned}
& \left.<0, \text { EVAL.c, n.s, v, G[n=AP } n_{1} n_{2}\right], E, D> \\
& \Rightarrow<0, \text { UNWIND.(), } n .(), v, G\left[n=A P n_{1} n_{2}\right], E,(c, s) . D> \\
& <0, E V A L . c, n . s, v, G[n=I N T i], E, D> \\
& \quad \Rightarrow<0, c, n . s, v, G[n=I N T i], E, D> \\
& \\
& \text { similarly for nodes BOOL b, NII, CONS } n_{1} n_{2} \text { and FUN f. }
\end{aligned}
$$

```
<o, UNWIND.(), n.s, v, G[n=AP \(\left.n_{1} n_{2}\right], E, D>\)
        \(\Rightarrow \ll 0\), UNWIND.(), \(n_{1} . n . s, v, G\left[n=A P n_{1} n_{2}\right], E, D>\)
\(<0\), UNWIND.(), \(n_{0} \cdot n_{1} \ldots n_{k} \cdot S, v, G\left[n_{0}=F U N f, n_{1}=A P n_{0} n_{1}{ }^{\prime}, \ldots, n_{k}=A P n_{(k-1)} n_{k}{ }^{\prime}\right], E[f=(k, c)], D>\)
        \(\Rightarrow \ll 0, c, n_{1}{ }^{\prime} \ldots n_{k}{ }^{\prime} . s, v, G\left[n_{0}=F U N A, n_{1}=A P n_{0} n_{1}{ }^{\prime}, \ldots, n_{k}=A P n_{(k-1)} n_{k}\right], E[f=(k, c)]\), \(>\)
\(<0\), UNWIND.(), \(n_{0} \cdot n_{1} \ldots n_{k} \cdot(), v, G\left[n_{0}=F U N f\right], E[f=(a, c)],\left(c^{\prime}, s^{\prime}\right) . D>\) and \(k<a\)
        \(\Rightarrow \ll 0, c^{\prime}, n k . s^{\prime}, v, G\left[n_{0}=F U N f\right], E[f=(a, c)], D>\)
<o, RET m.c, \(n_{1} \ldots n_{m} . n .(), v, G[n=\mid N T i], E,\left(c^{\prime}, s^{\prime}\right) . D>\)
        => <o, c', n.s', v, G[n=INT i], E, D>,
        similarly for nodes BOOL \(b, N I I\) and CONS \(n_{1} n_{2}\)
\(<0\), RET m.c, \(n_{1} \ldots n_{m} \cdot n . s, v, G\left[n=A P n_{1} n_{2}\right], E, D>\)
        \(\Rightarrow<0\), UNWIND.(), n.s, v, \(G\left[n=A P n_{1} n_{2}\right], E, D>\),
        similarly for \(n=\) FUN \(f\)
\(<0\), JFALSE I.c, s, true .v, G, E, D> \(\quad=><0, c, s, G, E, D>\)
\(<0\), JFALSE I.c, s, false.v, G, E, D> \(\quad \Rightarrow \ll 0, J M P\) I.c, s, G, E, D>
\(<0\), JMP I...LABEL I.c, s, G, E, D> \(\quad \Rightarrow \ll 0, c, s, G, E, D>\)
\(<0\), LABEL I.c, s, G, E, D> \(\quad \Rightarrow<0, c, s, G, E, D>\)
```


## Initial and final state of the G-machine

The initial configuration of the machine for the evaluation of the expression $e_{0}$ in environment $E_{0}$ is:


```
where c0 = E[eo)\ro;PRINT
```

This is a machine with an empty output, the code for evaluating the start expression ( $\mathrm{E}\left[\mathrm{e}_{0}\right] \mathrm{ro}_{0} 0$ ), an empty basic value stack, an empty graph, environment $\mathrm{E}_{0}$ containing the compiled code for the functions together with their arity, and an empty dump.

The machine stops when the state $<0,(0,(),(), G, E,()>$ has been reached.

## Differences between the abstract machines

Apart from the way in which the machines are described, they also differ as abstract graph rewrite machines. The most important distinctions are:

- The nodes in the graphs of both machines are different. The nodes in the ABC-machine contain an independent context and descriptor-id, whereas the nodes in the G-machine only hold a type identifier.
- The kind of graphs handled by the G-machine is different from the graphs treated in the ABC-machine. In the ABC-machine nodes have variable arity, while the nodes in the G-machine are either leave nodes or AP nodes with two arguments.

The graph Ackerman 23 is represented in the G-machine as AP (AP (FUN Ackerman) (INT 2)) (INT 3). Later on, variants of the G-machine are developed which use a more efficient representation for these graphs. [Burn 88, Johnson 87]

- The type list is the only data type available in the G-machine; special instructions are available to handle it. In the ABC-machine arbitrary constructors are handled.
As shown in the next chapter, the generated code for the machines also differs, e.g. the calling conventions of functions are different, and the B-stack is used in a more general way than the stack V .


## Differences between the descriptions

The G-machine is specified in another formalism than the ABC-machine, although the tuple describing the machine state shows some resemblance. The main differences between the descriptions are:

- The G-machine is not specified in a programming language; this implies that it can not be checked by a compiler, nor serve as a prototype implementation. The original specification of Johnson contains some errors (like unbound variables in the right hand side of a description rule), that would be spotted by a compiler.
- The G-machine description is not hierarchical, this implies that all items must be addressed at the same level. In the ABC-machine implementation details can be hidden in the micro-instructions.
- The instruction sequence executed is specified in the ABC-machine in a more direct way than in the G-machine. The flow of control is also more restricted in the G-machine, e.g. jumps are only forward and within the code resembling to one function. In later variants of the G -machine a more elaborated flow of control is defined [Johnson 87].
- The G-machine description is not complete, some cases (see EVAL and RET instructions) are specified by similarly for nodes.... This is not always perfectly clear (e.g. must the EVAL instruction for the node FUN if behave similar to the first rule, or to the second rule of the given specification).
- The described G-machine is essentially an interpreter, whereas the ABC-machine contains directly executable code. The instructions of the ABC-machine are functions that transform the machine state. The state transformations of the G-machine must be executed by some external function.


### 3.7 Discussion

This chapter presents a large and complex example of the description method introduced in the previous chapter. The abstract machine specification presented is actually the specification of the ABC-machine. It was not made after the
completion of the design, but the machine was actually designed by writing this specification. The description is used with success as prototype implementation and it appeared to be a useful definition for the actual implementation. It appeared to be easy to add some instructions and to experiment with slightly different instruction sets. The development of this specification was forcing a well structured design and the development time was spent on designing the machine instead of solving problems with its specification.

The ABC-machine is an imperative graph rewrite machine, it is an abstraction of concrete stack based machine architectures. Its instruction set provides an additional level of abstraction for the implementation of functional languages. The introduction of this level enhances portability, without a significant reduction in efficiency. The assembly level provided permits the use of labels and symbolic names. The execution speed of the prototype is sufficient for small test programs. The Miranda interpreter on a SUN 3/280 executes circa 10 instructions each second. The nfib-number for the ABC-code shown in section 5.6 is 1. However, the execution speed can increased almost an order of magnitude by omitting some of the code to give meaningful error messages in the microinstructions.

There are several implementations of the ABC-machine. The specification of the abstract ABC-machine appeared to be a very useful and a valuable definition of the abstract machine. The implementation in C of a parallel ABC-interpreter [Nöcker 89] follows the specification presented here very closely. The machine state is not passed from instruction to instruction, but as a global accessible data structure. By implementing the components of the ABC -machine in C a micro-instruction level is defined. The instructions are implemented using these micro-instructions. A state-of-the-art speed is achieved by the translation of ABC-code to MC68000 code for a SUN3/280 [Weijers 90, van Groningen 90].

# Chapter 4 Abstract Program Translation 

4.1 The description method
4.2 Bracket abstraction
4.3 The traditional description of bracket abstraction
4.4 Bracket abstraction described by syntactical matching
4.5 Bracket abstraction as abstract program transformation
4.6 Discussion

This chapter shows that functional programming languages can be used to describe program transformations and translations. The well-known bracket abstraction algorithm is used to compare this description method with some conventional descriptions. In the next chapter an elaborated example of the use of abstract program translations is shown.

The descriptions are called abstract program translations since the language to be transformed is represented by an abstract syntax tree (AST). In a functional language a data type is used to hold the AST. The structure of an AST is close to the structure of the language to be transformed. A translation is easier to express using the AST than using the syntax since the syntactical representation is just a linear sequence of characters, while an AST is a tree with the desired structure. Handling the syntactical structure in the description has as advantage that the syntactical sugar can be used in the description, but this description method appears to be only suited for small problems.

These data structures are not hidden in an abstract data type like the data structures used to represent the machine components in the previous chapters, but they are manipulated at the top level of the description. The translation functions use pattern matching on the data structure and deliver fragments of this data structure. Here the data structure is the representation of the objects manipulated, there is no reason at all to hide them in an abstract data structure. In the previous chapters the data structures themselves were irrelevant, just a carrier for the data was needed.

Efficiency and space consumption during the execution of the description is irrelevant. This implies that a simple data structure in close correspondence with the language to be transformed can be used. The data structures used in the description of program transformation do not resemble the data structure that would be used in a compiler. For a compiler efficiency and compactness of the data structures used are important constraints. This implies for instance the use of symbol tables.

In conventional descriptions of translations the language components handled are often surrounded by $\llbracket$ and $\rrbracket$ to distinguish them from the translation algorithm. In practice however, it appears to be impossible or unpractical to make a strict separation between translation algorithm and translated language.

The abstract program translation methodology will be explained in section 4.1. Then several descriptions of bracket abstraction will be given in order to make a comparison of the description methods possible. The specifications shown here are the original description (in section 4.3), the specification using the syntactical match (in 4.4) and the abstract program translation (4.5). In the next chapter an abstract program translation is used to describe the translation of Clean to ABC-code. In appendix A it is used to describe the mappping from ABC-assembler code to instructions. It is also used by the author to specify the translation from Miranda to Clean [Koopman 88].

### 4.1 The description method

In order to specify a translation a suited representation of the languages involved must be available. Here the languages involved are represented by a data type in a functional language. The defined data type to represent the syntax tree is called an abstract syntax tree (AST). The translation is described by a function that takes the AST to transform as argument and delivers the resulting AST.

A suited type must be defined for each language involved in the translation. The choice of a good definition of this type appears to be crucial to obtain an elegant specification. The type definition cannot be obtained in some mechanical way. Valuable guide-lines for this definition are: keep it small, simple and close to the syntax tree. A direct translation of the syntax to a data type is often a good starting point, redundant constructs (like infix notation) are removed unless it is essential to describe their conversion to function applications. It is generally more convenient to express additional information as annotations in the AST than to derive this information when it is needed. Fields to insert these annotations are added to the type as required.

An AST can be made almost generally applicable by inserting a variable field for the name of a node and a list of arguments. In this way a very general applicable expression tree is defined.

```
tree :"m Node name args
args == [tree]
name == [char]
```

Although convenient in some circumstances the general use of this kind of tree is not advocated. Due to the lack of distinguishing structure the type system does not help to check the construction of consistent trees. We advocate to use tailor made trees as shown below.

As first example, a suited AST for the $\lambda$-calculus [Church 32, Barendregt 84] is shown. The $\lambda$-calculus handles $\lambda$-terms; objects with an extremely simple structure. The syntax of $\lambda$-terms is

```
term := '(' term term ')' | an application
    ' }\lambda\mathrm{ ' var '.' term | an abstraction
    var a variable
```

A data structure suited to represent $\boldsymbol{\lambda}$-terms as AST in Miranda is

| term := | AP term term | 11 an application |
| :---: | :---: | :---: |
|  | ABS var term | \| |l an abstraction |
|  | VAR var | 11 a variable |
| var $==$ | [char] | II a possible repr |

Example: The $\lambda$-term $\lambda f .(\lambda g .(\lambda x .((f x)(g x))))$ is represented by the data structure
ABS " f " (ABS " g " (ABS " $x^{\prime \prime}$ (AP (AP (VAR "f") (VAR "x")) (AP (VAR " $g$ ") (VAR "x"))))
A variable occurs free in a term if it is not in the scope of an abstraction of that variable; it is bound otherwise. Using the tree representation of $\lambda$-expressions a variable is bound if it is a sub-expression, a leaf in the tree, of an abstraction of the same variable. The last occurrence of an abstraction of that variable seen from the root of the tree is said to bind the occurrence of that variable.

```
free_vars :: term -> [var]
lree_vars (AP t1 t2) = mkset (free_vars t1 ++ free_vars t2)
free_vars (ABS var term) = free_vars term -- [var]
free_vars (VAR var) = [var]
```

$m k s e t$ is the library function to remove duplicated elements from a list.
The replacement of a free variable by some term is called substitution. The process of (repeated) substitution is called reduction. An expression that can be reduced is called a redex (reducible expression). To keep the definition simple it is required that the expression substituted does not contain free variables that become erroneously bound through the substitution. The proper reduction of
such terms would require $\alpha$-conversion (see below), since our purpose is just to show the use of an AST to describe a language it is not treated here.

$$
\begin{aligned}
& x[x:=N] \equiv N ; \\
& y[x:=N] \equiv y, \text { if } x \neq y ; \\
&(\lambda x . M)[x:=N] \equiv \lambda x . M ; \\
&(\lambda y \cdot M)[x:=N] \equiv \lambda y \cdot(M[x:=N]) \text {, if } x \neq y ; y \text { is not allowed to occur free in } N \\
&\left(M_{1} M_{2}\right)[x:=N] \equiv\left(M_{1}[x:=N]\right)\left(M_{2}[x:=N]\right) ;
\end{aligned}
$$

```
substitute :: term -> var -> term -> term
```



A change of a bound variable to a fresh variable in a term is called $\alpha$-conversion.

```
\(\lambda x . M=\lambda y \cdot M[x:=y]\)
alpha_conversion :: var -> term -> term
alpha_conversion y (ABS \(\times \mathrm{m}\) ) \(=\) ABS y (substitute \(\mathrm{m} \times\) (VAR y\()\) )
```

$\beta$-reduction is defined by.

```
(\lambdax.N)M->\beta}N[x:=M
reduction :: term -> term
reduction (AP (ABS \times n)m}=\mathrm{ substitute n mm
```

This example shows that a data structure can be chosen in close correspondence with the syntactical structure. Some transformations are shown to be easily expressible.

### 4.2 Bracket abstraction

In order to compare the description method proposed here with some other description methods, the well-known bracket abstraction algorithm will be described in several formalisms. Bracket abstraction was introduced by Turner [Turner 79] as implementation methodology for functional languages. The theoretical background and soundness of abstraction can be found in combinatorial logic [Schönfinkel 24, Curry 58].

In order to use bracket abstraction as an implementation technique for functional programming languages, every function is treated as a higher order function. This implies that arguments are bound one by one in each function
application. Bracket abstraction is a program transformation that removes all variables (arguments) from a function body. Combinators from a small and fixed set are inserted in the function body to distribute the arguments during program execution. The absence of variables eliminates the need to maintain environments to associate actual arguments to variables during the execution of the program.

The language to transform is SASL [Turner 76], a simple predecessor of the language Miranda.

## Combinator reduction

An expression in SASL is reduced in two steps. First it is transformed to combinator code as shown below. Then the combinator expression is reduced by applying the reduction rules for these combinators. The reductions are carried out in the normal order; the leftmost redex in the expression is continuously rewritten. For some combinators it is necessary to evaluate one or more arguments before the rewrite rule can be applied.

### 4.3 The traditional description of bracket abstraction

The description given here is a recapitulation of the relevant parts of Turner's famous paper. Turner utilised an extended version of the description method used by Curry. In order to make this section self-contained no knowledge of Turner's article, nor of combinatorial logic is assumed.

## The basic abstraction algorithm

To remove variables from the source text the definition

$$
\operatorname{def} f x=E
$$

is considered. E is an expression in which function application is the only operation. It contains constants, including curried versions of operators, and variables. To transform an expression to its curried form functions are replaced by their curried equivalents and infix operators are replaced by the corresponding curried prefix operators. The example of the factorial function in the source language and in its curried version illustrates this transformation.

```
deffacn = n= 0-> 1; n*fac (n-1)
def fac n = cond (eq 0n) 1 (times n (fac (minus n 1)))
```

The variable x can be removed from the definition above by abstracting it from the expression. This is denoted by
def $f=[x] E$

In this definition, $[x] E$ denotes the result of the operation which removes all occurrences of $\times$ from $E$. It is pronounced 'abstract $x$ from $E$ '. The step taken above is correct when the law of abstraction holds. This law states that abstraction is the inverse of application.

$$
([x] E) x \quad=E
$$

A combinator is a function without free variables. To distribute the arguments in the expression the combinators $\mathbf{S}, \mathbf{K}$ and $\mathbf{I}$ are defined by their reduction rules.

| $\operatorname{Sfgx}$ | $=f x(g x)$ |
| ---: | :--- |
| $\mathbf{K x y}$ | $=x$ |
| IX | $=x$ |

Function application is denoted by juxtaposition, it associates to the left. Redundant brackets are usually omitted. Indicating the associations in the first combinator definition explicitly, it is written as:

$$
((S f) g) x=(f x)(g x)
$$

This implies that

```
deffxy = E must be read as
def (fx)y = E applying the abstraction rule yields
def fx = [y]E which is
def f}=[\textrm{x}]([y]E
```

The abstraction algorithm is best explained by viewing expressions as trees. The abstraction algorithm considers various cases of these trees:

- Abstraction of a variable over an application, a fork in the tree, yields a $\mathbf{S}$-combinator to distribute the argument over both sub-trees. The abstraction algorithm is recursively applied to both sub-trees.
- The abstraction of a variable from the same variable yields an I-combinator; the argument is needed here in the tree.
- The abstraction of a variable from all other variables and constants yields a K -combinator; the argument is not needed in this leaf of the tree and it is thrown away.
The abstraction algorithm is specified as ( $E_{1}$ and $E_{2}$ are arbitrary expressions)

| $[x]\left(E_{1} E_{2}\right)$ | $\Rightarrow \mathbf{S}\left([x] E_{1}\right)\left([x] E_{2}\right)$ |
| :--- | :--- |
| $[x] x$ | $\Rightarrow \mathbf{I}$ |
| $[x] y \quad$ | $\Rightarrow \mathbf{K} y \quad$, where $y$ is a constant or a variable other than $x$. |

## Proof:

Using the law of abstraction and the reduction rules for the combinators (the expression rewritten is underlined).

$$
\begin{aligned}
(\underline{[x] x}) x & =I x=x \\
(\underline{[x] y}) x & =\frac{1 x}{K} x=y \\
\left(\underline{[x]\left(E_{1} E_{2}\right)}\right) x & =\frac{\left.S\left([x] E_{1}\right)\left([x] E_{2}\right) x\right)}{\left(\left([x] E_{1}\right) x\right)\left(\left([x] E_{2}\right) x\right)} \\
& =\left(E_{1} E_{2}\right) \text { by induction to the structure of terms. }
\end{aligned}
$$

The abstraction of the argument from the definition for the factorial function in curried form

```
def fac n = cond (eq 0 n) 1 (times n (fac (minus n 1)))
```

yields

$$
\begin{aligned}
& \text { def fac }=\mathbf{S}(\mathbf{S}(\mathbf{S}(\mathbf{K} \text { cond) ( } \mathbf{S}(\mathbf{S}(\mathrm{K} \mathrm{eq})(\mathbf{K} 0)) \mathbf{I})(\mathrm{K} \text { 1) } \\
& \text { ( } \mathbf{S} \text { (S (K times) I) (S (K fac) (S (S (K minus) I) (K 1) }) \text { ) }
\end{aligned}
$$

## Improving the generated combinator code

The definitions given above form a complete algorithm to remove variables from an expression. However, the combinator expression obtained by this algorithm is rather long winded. This is due to the nature of the abstraction algorithm used; each argument is brought down to every leaf of the expression tree. When the argument is not needed in a leaf it is discarded there by a K -combinator. The size of the expression can be reduced by discarding arguments as soon as possible and directing them only to the branches of the expression tree where they are needed. This is done by introducing some additional combinators and applying improvement rules from the leaves to the top of the tree ${ }^{1}$.

The combinators needed to improve the code are defined by the rules:
$B f g x=f(g x)$
Cfgx $=f x g$
The rules to improve the code are:

1. The distribution (by the S-combinator) of an argument over two sub-trees is not needed when the argument is discarded right away in both sub-trees (by a K-combinator). The argument can be discarded immediately.

[^4]2. The distribution of an argument over two sub-trees is not needed if it is discarded in the left sub-tree and becomes the right sub-tree by an I-combinator.
3. The distribution of an argument which is thrown away in the left branch is equivalent with sending the argument to the right tree by a B-combinator.
4. The same holds when the argument is not needed in the right sub-tree. The argument can be passed to the left branch by a C-combinator.
These optimization rules are expressed formally as

| $\mathbf{S}\left(\mathbf{K} \mathrm{E}_{1}\right)\left(\mathbf{K} \mathrm{E}_{2}\right)$ | $\Rightarrow \mathbf{K}\left(\mathrm{E}_{1} \mathrm{E}_{2}\right)$ |  | II improvement rule 1 |
| :---: | :---: | :---: | :---: |
| $\mathbf{S}\left(\mathrm{KE} \mathrm{E}_{1}\right) \mathrm{I}$ | $\Rightarrow \mathrm{E}_{1}$ |  | 1 improvement rule 2 |
| $S\left(K E_{1}\right) E_{2}$ | $\Rightarrow B E_{1} E_{2}$ | , if no earlier rule applies | 11 improvement rule 3 |
| $S \mathrm{E}_{1} \quad\left(\mathrm{~K} \mathrm{E}_{2}\right)$ | $\Rightarrow C E_{1} E_{2}$ | , if no earlier rule applies | $1 /$ improvement rule 4 |

The first rule states that $\left.\mathbf{S}\left(K E_{1}\right)\left(K E_{2}\right)\right)$ and $K\left(E_{1} E_{2}\right)$ are equivalent expressions. Such a relation can be proven easily. An arbitrary argument, $x$, is supplied and the reduction rules for the combinators are used to transform the expression.

```
Proof:
- rule 1
\(S\left(K E_{1}\right)\left(K E_{2}\right) x=\{\) rewrite rule for \(S\) \}
\(\overline{\left(K E_{1} x\right)}\left(K E_{2} x\right) \quad=\) (rewrite rule for \(K\) \}
\(E_{1}\left(K E_{2} x\right)=\{\) rewrite rule for \(K\) \}
    \(\mathrm{E}_{1} \mathrm{E}_{2}\)
and
    \(\mathbf{K}\left(E_{1} E_{2}\right) \times \quad=\{\) rewrite rule for \(K\}\)
\(\mathrm{E}_{1} \mathrm{E}_{2}\)
hence \(S\left(K E_{1}\right)\left(K E_{2}\right)\) and \(K\left(E_{1} E_{2}\right)\) are extensionally equal.
- rule 2
    \(\underline{s\left(K E_{1}\right) \mid x}=\left(K E_{1} x\right)(1 x)=E_{1}(\underline{I x})=E_{1} x\)
hence \(\mathbf{S}\left(K E_{1}\right) \boldsymbol{I}\) and \(E_{1}\) are extensionally equal.
- rule 3
    \(S\left(K E_{1}\right) E_{2} x \quad=\left(K E_{1} x\right)\left(E_{2} x\right)=E_{1}\left(E_{2} x\right)\)
\(B E_{1} E_{2} x \quad=\overline{E_{1}\left(E_{2} x\right)}\)
hence \(\mathbf{S}\left(\mathrm{KE}_{1}\right) \mathrm{E}_{2}\) and \(\mathbf{B} \mathrm{E}_{1} \mathrm{E}_{2}\) are extensionally equal.
- rule 4
    \(S E_{1}\left(K E_{2}\right) \times=E_{1} \times\left(K E_{2} \times\right)=E_{1} \times E_{2}\)
\(C E_{1} E_{2} x=E_{1} \times E_{2}\)
```

hence $S E_{1}\left(K E_{2}\right)$ and $C E_{1} E_{2}$ are extensionally equal.

This proof shows also that the improved expressions require less reduction steps; the expressions are indeed enhanced.

Using these rules the combinator code for the factorial function becomes:

```
def fac = S (C (B cond) (eq 0)) 1(S tlmes (B fac (C minus 1)))
```

These simple optimizations give a considerable reduction of the size of the expressions. Reductions of $50 \%$ and more are common. The code can be further improved by defining additional combinators and considering more special cases [Turner 79, Peyton Jones 87], since our aim is not to describe the best possible fixed combinator generation this is not done.

## Data structures

There is only one data structure defined in SASL; the list. Lists are built by the pairing operator and the empty list nII. The pairing operator (called cons) is represented by an infix colon ' $\because$ '. In prefix form it is the combinator $P$ (for 'pair'). The combinator $U$ (for 'uncurry'2) unpacks pairs.

$$
U f(P x y)=f x y
$$

The left hand side of a function definition in SASL may be, or contain, a list. Hence the abstraction algorithm has to be extended with respect to abstractions of pairs. The Pair is unpacked and both elements are abstracted.

$$
[\mathbf{P} \times \mathrm{y}] \mathbf{E} \quad \Rightarrow \mathbf{U}([\mathrm{x}]([y] \mathrm{E}))
$$

The use of this abstraction rule is illustrated by the example shown below.

```
def f(a:b) = a +b becomes in prefix form
deff(Pab) = plus ab abstracting the function argument
def f = [Pab] (plus a b) which is after abstraction
def f = S (B plus (U K)) (U (K I))
```


## Local definitions

In SASL, every expression can be accompanied by a sequence of locally defined functions. This is transformed to a single combinator expression in three steps.

1. The arguments of the locally defined functions are removed.

$$
\begin{aligned}
E_{1} \quad \text { where } f x & =E_{2} \\
g y & =E_{3} \\
E_{1} \quad \text { where } f & =[x] E_{2} \\
g & =[y] E_{3}
\end{aligned} \quad \text { is in the first step transformed to: }
$$

[^5]2. The sequence of locally defined functions is grouped to a single local definition. This is done by packing them in a list.

| $E_{1} \quad$ where $f$ | $=E_{2}$ |
| ---: | :--- |
| $g$ | $=E_{3}$ |
| $E_{1}$ where $(f: g)$ | $=\left(E_{2}: E_{3}\right)$ |$\quad$ is in the second step transformed to:

3. Finally, the resulting local definition is removed. This is done by abstracting the list containing the locally defined function names from the expression $E_{1}$. The locally defined function body is supplied as the corresponding argument.
$E_{1}$ where $f=E_{2}$ is transformed to:
([f] $\left.E_{1}\right) E_{2}$

## Recursion

There is one further complication to handle; when the locally defined functions are (mutual) recursive, the algorithm outlined above would not abstract the names of the locally defined functions from their body. This is illustrated by an example of the code produced by the algorithm above.
fac 3 where fac $n=n=0 \rightarrow 1 ; n$ * fac $(n-1)$
(C I 3) (S (C (B cond) (eq 0)) 1 (S times (B fac (C minus 1$)$ ))
The code produced is obviously wrong! It still contains the identifier fac. Recursion can be handled correctly by using the fixed point combinator Y. A term $e$ is a fixed point of a function $f$ when $f e=e$. The fixed point combinator is defined such that, for any function $f, Y f$ is a fixed point of $f$.

$$
Y f=f(Y f)
$$

A correct algorithm to handle local definitions removes the occurrence of the locally defined function name from its body. Then, this function is removed from the expression

| $E_{1} \quad$ where $f=E_{2}$ | is first transformed to |  |
| :--- | :--- | :--- |
| $E_{1} \quad$ where $f=Y([f]$ | $\left.E_{2}\right)$ | this is transformed to |
| $\left([f] E_{1}\right)\left(Y\left(f f E_{2}\right)\right)$ |  |  |

Using this transformation rule, the factorial example above becomes
(C 13$)(\mathrm{Y}(\mathrm{B}(\mathrm{S}(\mathrm{C}(\mathrm{B}$ cond (eq 0))1))(B(S times)(C(BSK)(C minus 1))))

Global function identifiers are not removed from the expression at compile time, but they are replaced by the associated combinator expression when they are encountered at run-time. So, recursive functions are no problem. Since this implementation is outside the world of fixed combinators discussed here, it is not treated in the sequel.

## Summary of the abstraction algorithm

In order to make a comparison of the description methods, the rules of the abstraction algorithm are summarised. The equivalent of these rules in the other formalisms will be given below.

The combinators used in the abstraction process are defined by their reduction rules.

$$
\begin{array}{ll}
S f g x & =f x(g x) \\
K x y & =x \\
I x & =x \\
B f g x & =f(g x) \\
C f g x & =f x g \\
U f(P x y) & =f x y \\
Y f & =f(Y f)
\end{array}
$$

Arguments are removed from the definition ( $E$ is an expression in which function application is the only operation)

$$
\operatorname{def} f x=E
$$

by abstracting them from the body. This is denoted as

$$
\operatorname{def} f \quad=[x] E
$$

The abstraction rules must by applied in the order they are listed here.

$$
\begin{aligned}
& {[x]\left(E_{1} E_{2}\right) \Rightarrow \mathbf{S}\left([x] E_{1}\right)\left([x] E_{2}\right)} \\
& {[P \times y] E \quad \Rightarrow U([x]([y] E))} \\
& {[x] x \quad \Rightarrow I}
\end{aligned}
$$

$$
[x] y \quad \Rightarrow K y \quad \text {, where } y \text { is a constant or a variable other than } x \text {. }
$$

Local definitions must be removed from the expression before the abstraction is done. The algorithm to remove local definitions is explained by showing how two local definitions with one argument each, are removed.

$$
\begin{aligned}
E_{1} \text { where } f x & =E_{2} \\
g y & =E_{3}
\end{aligned} \quad \text { the arguments are removed first }
$$

```
\(\rightarrow E_{1}\) where \(f \quad=[x] E_{2}\)
    \(g=[y] E_{3}\)
\(\rightarrow \quad E_{1} \quad\) where ( \(f: g\) ) \(=\left([x] E_{2}:[y] E_{3}\right)\)
\(\rightarrow \quad E_{1} \quad\) where \(\left.(f: g)=Y\left([(f: g)]\left([x] E_{2}: l y\right] E_{3}\right)\right)\)
\(\rightarrow \quad\left([(f: g)] E_{1}\right)\left(Y\left([(f: g)]\left([x] E_{2}:[y] E_{3}\right)\right)\right)\)
```

then the definitions are paired next, the recursion is handled lastly the definition is removed

The improvement rules to reduce the size of an expression are

| $\mathbf{S}\left(K E_{1}\right)\left(K E_{2}\right)$ | $\Rightarrow K\left(E_{1} E_{2}\right)$ |  |
| :--- | :--- | :--- |
| $\mathbf{S}\left(K E_{1}\right) I$ | $\Rightarrow E_{1}$ |  |
| $\mathbf{S}\left(K E_{1}\right) E_{2}$ | $\Rightarrow B E_{1} E_{2}$ | , if no earlier rule applies |
| $\mathbf{S} E_{1}$ | $\left(K E_{2}\right)$ | $\Rightarrow C E_{1} E_{2}$ |

The reduction order used in the SASL system is not specified in this formalism.

## Suitability of the description method

Although this description method is syntactically appealing and very close to the notation of abstraction used in mathematical textbooks, it has a number of disadvantages. A serious drawback is that the various algorithms introduced have no intrinsic names. The absence of these names turns out to be problematic when it is not quite clear from the context which algorithm is meant and the absence of these names is especially annoying when several algorithms are merged. It is necessary to combine transformations to achieve a more efficient algorithm, or to handle more complex cases. Another limitation of the syntactical match is that it is only adequate to describe relatively simple transformations. For instance, in the original description the transformations necessary to remove locally defined functions are not defined by some general applicable transformation rule, but they are specified only by an example. The introduction of additional description tools is necessary to define such a transformation rule. Also the entire transformation from SASL to SK-combinator code is not specified as one algorithm, but as a number of transformations and some textual indication how they must be applied.

### 4.4 Bracket abstraction described by syntactical matching

The obvious way to handle the problem with the names of the transformations is to define the transformations as functions. This spoils the nice syntactical match a little bit, since there are now two kinds of objects; the syntactical objects to transform and the functions performing the transformations. An attempt to keep these matters separated is the use of Scott brackets, $\mathbb{L}$ and $\mathbb{I}$, to indicate the match on syntactical objects. These symbols are widely used in denotational semantics to indicate that syntactical objects are handled [Gordon 79, Schmidt 86, Fehr 89].

The combinator definitions and the algorithm for bracket abstraction shown below are similar to the definitions in textbooks on the implementation of functional programming languages [Peyton Jones 87]. The transformation shown here is extended to deal with SASL expressions instead of $\lambda$-terms. The algorithm is expressed in a kind of functional programming language. This language is usually not properly defined, but so simple that no formal definition is felt to be necessary. However, good text books on denotational semantics give a firm definition. The use of this description method as if it were a functional programming language suggests a certain accuracy in the description which is usually not there, e.g. a last alternative to handle the 'do nothing' case is typically written in half of the situations where it is required.

## Combinators used in the abstraction algorithm

The combinators are defined by their reduction rules.

| $S f g x$ | $\rightarrow f x(g x)$ |
| :--- | :--- |
| $K x y$ | $\rightarrow x$ |
| $i x$ | $\rightarrow x$ |
| $B f g x$ | $\rightarrow f(g x)$ |
| $C f g x$ | $\rightarrow f x g$ |
| $U f(P x y)$ | $\rightarrow f x y$ |
| $Y f$ | $\rightarrow f(Y f)$ |

## The expressions handled

The expressions handled are SASL expressions in applicative form. The keywords are printed in bold face. Combinators are printed in uppercase. All other expressions are variables: either SASL variables or transformation variables, they are printed in lower case.

The following conventions are used:

| $e, e_{1}, p, q$ | are arbitrary expressions; |
| :--- | :--- |
| $f, f_{i}$ | are function identifiers, or function identifiers with some formal arguments; |
| $x_{i} y, a_{1}$ | are arguments (variables or pairs); |
| $v$ | is a variable; |
| cv | is a constant or a variable; |
| def,def | are definitions; |
| defs,locals | are sequences of (local) definitions. |

By applying the transformations described here a SASL expression can be transformed to a pure combinator expression. $\mathbb{C} \mathbb{I}$ @ Compiles the expression e, i.e. transforms it to a combinator term. The compilation algorithm is described by a number of transformation rules. The first transformation rule alternative applicable must be used.


```
C[迹were locals \rrbracket = RIs e\llbracketlocals \rrbracket
C[c] = = C
```


## The abstraction algorithm

To avoid unnecessary scanning of the expression tree the optimization rule is invoked whenever appropriate during the abstraction．In the previous descrip－ tion this could not be indicated since the optimization rule has no name in that formalism．
$A \times \llbracket e \rrbracket A b s t r a c t s$ the argument x from the expression e.

```
\(A \times \llbracket\) e where locals \(\rrbracket=A \times \llbracket\) Ris \(e \llbracket\) locals \(\rrbracket \rrbracket\)
\(A \times \llbracket e_{1} e_{2} \rrbracket=O p t \mathbb{I} S\left(A \times \llbracket e_{1} \mathbb{I}\right)\left(A \times \llbracket e_{2} \mathbb{I}\right) \rrbracket\)
\(A(P \times y) \llbracket \vee \mathbb{Z} \quad=U(A \times(A y \llbracket \vee \mathbb{I}))\)
\(A \times \llbracket x \rrbracket \quad=1\)
\(\mathbf{A} \times \llbracket \mathrm{Cv} \mathbb{Z}=\mathrm{Kcv}\)
```

Aa【 $\ddagger \mathbb{D}$ Astracts all arguments from a function definition．

$$
\begin{array}{ll}
\text { Aa【 P } a_{1} a_{2}=e \mathbb{I} & =P a_{1} a_{2}=e \\
\text { Aa【I } f=e \mathbb{I} & =\operatorname{Aa\llbracket } f=A a \llbracket e \mathbb{I}] \\
\text { Aa【def } \mathbb{d} & =\operatorname{def}
\end{array}
$$

Rls e【locals $\rrbracket$ Removes the local definitions from the expression e．Arguments of the definitions are removed and they are grouped to a single definition．Then this definition is removed．

Ris $\mathrm{e} \llbracket$ locals $\mathbb{\rrbracket} \quad=\quad$ RIe $\mathbb{C}$ Col【 Aal locals $\mathbb{\rrbracket} \mathbb{\rrbracket}$
RI e IIlocals 】 Removes a single local definition from the expression e．Recursion in the definition is removed by a $Y$ combinator and abstracting the name of the function from its body．

$$
\operatorname{RL} \theta_{1} \llbracket f=e_{2} \mathbb{\rrbracket} \quad=\left(A f \mathbb{C} \theta_{1} \mathbb{I}\right)\left(\mathcal{Y}\left(A f \mathbb{C} e_{2} \mathbb{I}\right)\right)
$$

Aal【 defs $\mathbb{\rrbracket}$ Abstract the arguments from each definition in a list of definitions．

$$
\text { Aal【⿺ } \operatorname{def}_{1} \ldots \text { def }_{n} \mathbb{I} \quad=\text { Aa【 } \operatorname{def}_{1} \mathbb{\rrbracket} \ldots \text { Aal【 } \operatorname{def}_{n} \mathbb{\rrbracket}
$$

Col $\mathbb{I}$ defs $\mathbb{I}$ Collects a list of local definitions to a single definition．Definitions are paired from back to front．

```
ColII \(f_{1}=e_{1}\)
    \(t_{n-1}=e_{n-1}\)
    \(f_{n}=e_{n} \quad \mathbb{D}=\operatorname{Col} \mathbb{L} \quad f_{1}=e_{1}\)
    \(P f_{n-1} f_{n}=P e_{n-1} e_{n} \rrbracket\)
Col II def \(\mathbb{I}=\operatorname{def}\)
```


## Improving the generated combinator code

Opt［I e $\mathbb{\square}$ Optimizes the expression e．It is invoked by the abstraction rule．So，no scanning of the expression tree is needed．

$$
\begin{aligned}
& \text { Opt【S } S\left(K e_{1}\right)\left(K e_{2}\right) \rrbracket=K\left(e_{1} e_{2}\right) \\
& \text { Opt[IS } \left.\left(K e_{1}\right) 1\right]=e_{1} \\
& \text { Opt【S } S\left(K e_{1}\right) e_{2} \rrbracket=B e_{1} e_{2} \\
& \text { Opt【S } e_{1}\left(K e_{2}\right) \mathbb{I}=C e_{1} e_{2} \\
& \text { Opt【e] }=e
\end{aligned}
$$

## Reduction

Reduction is specified by the semantics of the combinator expressions．The order in which the rewrite rules are applied to the various parts of the expres－ sion is determined by the normal order reduction strategy．This means that the leftmost redex is rewritten．Sometimes it is necessary to reduce some sub－ expressions before the topmost redex term can be rewritten．For instance， before an addition can be performed both arguments must be reduced．Two functions are used to specify the reduction process．Reduce determines the order of reductions．First the left sub－tree of an application node is reduced．Then Rewrite is applied to the top node．This function tries to apply the reduction rules for the combinators，after a rewriting the reduction is continued when appropriate．These functions will transform an expression until the top cannot be rewritten any more．Then the expression is said to be in head normal form， hnf．

| RedII $e_{1} e_{2} \mathbb{I}$ Red［I］ 1 | $\begin{aligned} & =\operatorname{Rew} \mathbb{I} \operatorname{Red}\left[e_{1} \mathbb{I} e_{2} \mathbb{\rrbracket}\right. \\ & =e \end{aligned}$ |
| :---: | :---: |
| Rew［ $S^{\text {e }} \mathrm{e}_{1} \mathrm{e}_{2} \mathrm{e}_{3} \mathbb{\square}$ | $=\operatorname{Red}\left[e_{1} e_{3}\left(e_{2} e_{3}\right) \mathbb{l}\right.$ |
| Hew［IK $\mathrm{e}_{1} \mathrm{e}_{2} \mathrm{I}$ | $=\operatorname{RadI} \mathrm{e}_{1} \mathbb{I}$ |
| Rewllel | ＝Redll $\mathrm{e}^{\text {I }}$ |
| Rew［ $\mathrm{Bl}_{\left.\mathrm{e}_{1} \mathrm{e}_{2} \mathrm{e}_{3}\right]}$ | $=\operatorname{Red}\left[e_{1}\left(e_{2} e_{3}\right)\right]$ |
| Rew［LC $\mathrm{el}_{1} \mathrm{e}_{2} \mathrm{e}_{3} \mathbb{I}$ | $=\operatorname{Red}\left[e_{1} e_{3} e_{2} \rrbracket\right.$ |
| Rew［U $\backslash \mathrm{e}_{1} \mathrm{e}_{2}$ ］ | ＝U $\mathrm{e}_{1}$ Red［ $\mathrm{e}_{2} \rrbracket$ |
| RewII EQ $\mathrm{e}_{1} \mathrm{e}_{2} \mathrm{I}$ | $=\operatorname{Red}\left[\theta_{1} \mathbb{I}=\right.$ Red $\left[e_{2} \rrbracket\right.$ |
| Rew［IPLUS $\mathrm{e}_{1} \mathrm{e}_{2} \mathbb{\square}$ | $=\operatorname{Red}\left[\mathbf{e}_{1} \mathbb{\square}+\right.$ Red $\mathbb{C} \mathbf{e}_{\mathbf{2}} \mathbb{\square}$ |


| Rew［L MINUS $\mathrm{e}_{1} \mathrm{e}_{2} \rrbracket$ | $=$ Red［⿺𠃊 $e_{1} \mathbb{\square}$－Red $\mathbb{e} e_{2} \mathbb{\square}$ |
| :---: | :---: |
| Rew［ITMES $\mathrm{e}_{1} \mathrm{e}_{2}$ 】 | $=\operatorname{Red}\left[e_{1} \rrbracket \times R e d\left[e_{2} \mathbb{I}\right.\right.$ |
| RewIL COND $e_{1} e_{2} e_{3} \mathbb{\square}$ | $=\mathrm{IF}$ RedII $\mathrm{e}_{1} \mathbb{I} \mathrm{e}_{2} \mathrm{e}_{3}$ |
| Rew［L］ | $=0$ |
| $U e_{1} \llbracket P e_{2} e_{3} \rrbracket$ | $=\operatorname{Red} \mathbb{L} \mathbf{e}_{1} \mathbf{e}_{2} \mathrm{e}_{3} \mathbb{I}$ |
| If TRUE $\mathbf{e}_{1} \mathbf{e}_{2}$ | $=\operatorname{Red}\left[\mathrm{e}_{1}\right]$ |
| IF FALSE $\mathbf{e}_{1} \mathbf{e}_{\mathbf{2}}$ | $=\operatorname{Red}\left[\mathrm{e}_{2} \mathrm{I}\right.$ |

The strategy and the reduction rules are combined into a single algorithm to obtain a feasible implementation．This is similar to the combination of the func－ tional reduction strategy and the rewrite rules for Clean described in the next chapter．It is possible to specify these algorithms separately，but that involves a lot of overhead to scan the tree for the next redex and some additional construc－ tors to pass information from the rewrite algorithm to the strategy．This would resemble very much to the two level approach proposed by van Eekelen and Plasmeijer［van Eekelen 86，88］．

## Suitability of the description method

Unfortunately there is no strict separation between syntactical objects and the transformation functions in the description above．In the rules for abstraction， algorithmical components are used inside the syntactical region without being syntax．The expressions must be evaluated to become syntactical objects． Sometimes the symbol $\equiv$ is used instead of $=$ to indicate that the right hand side of the rule can be used as a syntactical object．In the Abstract arguments rule，the arguments，$a_{1}$ ，are taken from the syntactical region to the transformation region without notice．Also matching arbitrary variables and specific variables is somewhat tricky．The difference between $e_{1}, e_{2}$ and $x$ and $v$ in the abstraction rule is not very clear．The fact that $v$ is a variable in the rule to abstract a pair is quite subtle．In the abstraction rule there is even a match on syntactical structure outside the brackets．Sometimes uppercase characters are used for constants and lower case characters for variables，but this cannot be used consistently．The syntactical match can be used for a single context free transformations of lan－ guages with a relative simple structure，but for richer languages and more complex transformations it becomes very tricky．For instance，in the translation of Miranda to Clean it is necessary to distinguish between an arbitrary argu－ ment，an argument that is a variable，an argument that is a variable equal to an other argument，an argument that is an arbitrary constructor with a variable number of sub－arguments，an argument that is a member of a specific class of predefined values and an argument that is a specifc symbol［Koopman 88］．

### 4.5 Bracket abstraction as abstract program transformation

Since it appears to be impossible to keep the objects handled and the algorithm separated we propose to express both in a functional programming language. The essential differences between the description employing syntactical matching and the description in a functional language is, that the language constructs to be transformed are represented as syntactical objects between $\mathbb{I}$ and $\mathbb{\rrbracket}$ in the description above, while they are represented by a data structure in the functional description.

## The expressions handled

The SASL expressions are represented by an AST. The Miranda data structure expr is used to hold it.
expr
$::=$ AP expr expr
WHERE expr [def]
VAR var
INT num
BOOL bool
S |K |I |B |C
Y
P |U |NIL
COND
EQ | PLUS | MINUS |TIMES

A definition has a left-hand side, lhs, that is either a variable with a sequence of arguments or a pair. An argument is either a variable or a pair of arguments. In order to avoid the construction of an additional data-type it is represented by an expression. The right hand side, rhs, of a definition may be any expression.

```
def :== DEF lhs rhs
lhs == expr |l A lhs is a pair or a var with arguments.
rhs == expr |l All expressions are allowed on the ms
var == [char]
```


## Combinators used in the abstraction algorithm

The rewrite behaviour of the combinators defined is specified by the function rewrite. The combinators themselves can not be functions since trees which are not in head normal form must be handled (e.g. by the optimization rules). So, an interpreting function is used.

```
rewrite :: expr -> expr
rewrite (AP (AP (AP S f) g) \(x\) ) \(=A P(A P f x)(A P g x)\)
rewrite (AP (AP K x) y) = \(\quad\) )
rewrite (AP|x) \(=x\)
rewrite (AP Y f) \(\quad=A P f(A P Y f)\)
rewrite (AP (AP (AP B f) g) \(x\) ) \(=A P f(A P g x)\)
rewrite (AP (AP (AP C f) g) \(x\) ) \(=A P(A P f x) g\)
rewrite (AP (AP Uf) expr) = unpair \(f\) (reduce expr)
rewrite expr \(=\) expr
unpair \((\) (AP (AP Px)y)) \(=A P(A P f x) y\)
```

The rule for $U$ is special since it requires the reduction of a sub-term before it can be applied. Basic operations like PLUS, MINUS, TIMES, EQ and COND, also require reduction before they can be applied. The function reduce evaluates an expression tree to head normal form. It is defined below.

## The abstraction algorithm

By applying the transformation rules, an expression representing SASL program is transformed to a pure combinator expression. This process is initiated by the function compile.

```
compile :: expr -> expr
compile (AP e1 e2) = AP (compile e1) (compile e2)
compile (WHERE e defs) = optimize (remove_where e defs)
compile expr = expr
```

The abstraction algorithm expressed in Miranda reads:

```
abstract :: expr -> expr -> expr
abstract x (WHERE e defs) = abstract x (remove_where e defs)
abstract x (AP e1 e2) = AP(AP S (abstract x e1)) (abstract x e2)
abstract (AP (AP P x y)(VAR z) = AP U (abstract x (abstract y (VAR z)))
abstract var exp = I , if var = exp
= AP K exp , otherwise
```

Locally defined functions are removed from an expression as outlined above.

```
remove_where :: expr -> [def] -> expr
remove_where expr defs
= AP (abstract f expr) (AP Y (abstract f body))
    where DEF f body = collect (map abs_args defs)
```

Arguments are removed from a function definition by abstracting them from the function body. As a special case the left hand side of a definition can also be a pair.

```
abs_args :: del -> def
abs_args (DEF (AP (AP P h) t) body) = DEF (AP (AP P h) t) body
abs_args (DEF (AP |x) body) = abs_args (DEF f (abstract x body))
abs_args (DEF f body) = DEF fbody
```

The sequence of local function definitions is combined to one definition by pairing them from back to front.

```
collect :: [def] -> def
collect (det:defs)
= def , if defs = []
= DEF f body ,otherwise
    where f = AP (AP P f1)t2
        body = AP (AP P body1) body2
        DEF f1 body1 = def
        DEF f2 body2 = collect defs
```

The resulting AST fits nicely in the data structure expr, but variables and local definitions do not occur.

## Improving the generated combinator code

The function optimize applies the improve rules to every application, starting in the leaves of the expression tree. The optimization rules are not included in the abstraction algorithm to enable the observation of the effect of these improvements.

```
optimize :: expr -> expr
optimize (AP e1 e2) = improve (AP (optimize e1) (optimize e2))
optimize expr = expr
improve :: expr -> expr
improve (AP (AP S (AP K e1)) (AP K e2)) = AP K (improve (AP e1 e2))
improve (AP (AP S (AP K e1)) I) = e1
improve (AP (AP S (AP K e1)) e2) = AP (AP B e1) e2
improve (AP (AP S e1) (AP K e2)) = AP (AP C e1) e2
improve expr = expr
```


## Executing the generated combinator code

By applying the given rewrite rules the value of the resulting combinator expression can be computed. Two functions are used to describe the reduction to
hnf. Reduce controls the global order of reductions and rewrite performs the reductions, after a rewrite the reduction is continued when appropriate. Rewrite invokes the required reductions for rewrite rules that require arguments in hnf, a function corresponding to such a top node handles the reduction according to that node.

```
reduce :: expr -> expr
reduce (AP e1 e2) = rewrite (AP (reduce e1) e2)
reduce exp = exp
```

```
rewrite :: expr -> expr
rewrite (AP (AP (AP S f)g)x) = reduce (AP (AP fx) (AP gx))
rewrite (AP (AP K x) y) = reduce }
rewrite (AP | X) = reduce }
rewrite (AP Y f) = reduce (AP f (AP Y f))
rewrite (AP (AP (AP B f)g)x) = reduce (AP f (AP gx))
rewrite (AP (AP (AP C f) g) x) = reduce (AP (AP fx)g)
rewrite (AP (AP U f) exp) = unpair f (reduce exp))
rewrite (AP (AP PLUS x) y) = plus (reduce x) (reduce y)
rewrite (AP (AP MINUS x) y) = minus (reduce x) (reduce y)
rewrite (AP (AP TIMES x) y) = times (reduce x) (reduce y)
rewrite (AP (AP EQ x) y) = eq (reduce x) (reduce y)
rewrite (AP (AP (AP COND c) t) e = cond (reduce c) te
rewrite exp = exp
unpair f (AP (APPx)y) = reduce (AP (AP fx) y)
plus (INT i1) (INT i2) = INT (i1 + i2)
minus (INT i1) (INT i2) = INT (i1-i2)
times (INT i1) (INT i2) = INT (i1 * i2)
eq e1 e2 = BOOL (e1 = e2)
cond (BOOL True) t e = reduce t
cond (BOOL False) t e = reduce e
```

The reduction rules defined here perform a correct string reduction of the expression. A more efficient implementation of the reduction of these expressions will use graph reduction. Using graph reduction every sub-expression is reduced at most once.

## Suitability of the description method

The overhead involved in implementing the abstraction algorithm in a functional language appears to be very limited. This implementation effort is rewarded by a very clear and executable specification that can be partially checked by a Miranda implementation. The tree-like data structure needed to represent the SASL expressions appears to be essential in the abstraction process.

### 4.6 Discussion

In this chapter three descriptions of the language transformation known as bracket abstraction are compared. The two conventional descriptions use special purpose notations for the descriptions and are hence syntactically appealing. A drawback of the first description is that it lacks intrinsic names for the individual transformations. Only a set of transformation rules is given, but the combination of these rules to a complete algorithm cannot be described in this formalism. The second description possesses names for the individual transformations, the language to transformed is distinguished from the description language by surrounding it with special symbols. Unfortunately, it is generally not possible to make such a separation between both languages using this notation.

Using a programming language for the specification has the advantages, as said before, that the description can be partially checked by the implementation of the description language and that it is executable. This is a valuable test for the combinator code generated; the code generated is quite unreadable for human beings. Also the effect of the improvement rules can be determined easily using an executable description; apply the function optimize to some (generated) expressions. A possible drawback of such a description could be the implementation overhead involved, since no special purpose syntax can be introduced the description might become very long. But, we have shown that this overhead is very small when the functional programming language Miranda is used as description formalism. The obtained description is at least as clear as the mathematical descriptions.

In conclusion it can be stated that functional descriptions are very convenient to express these kind of program transformations. Descriptions given in this way tend to be correct, complete, compact, clear and executable. In the next chapter an abstract program transformation will be given for a more complex translation.

# Chapter 5 Translating Clean to ABC-Code 

### 5.1 Graph rewriting in Clean

5.2 Graph rewriting on the ABC-machine
5.3 The AST for Clean
5.4 Compilation
5.5 The run-time system
5.6 Curried functions
5.7 Optimizations
5.8 Discussion

In this chapter a translation of Clean to ABC -code is treated as a real life example of the abstract program transformation method introduced in the previous chapter. This transformation is rather complex since the relatively high-level declarative language Clean is transformed into a low-level imperative assembly language.


Fig 5.1 Transformation of Clean to ABC-code as a phase in the compilation process.
The goal of the translation is, of course, to transform Clean programs to equivalent and efficient ABC-programs. Since ABC-code is translated to (or interpreted by) various concrete machines it is not possible to assign execution times
to the instructions, hence it is impossible to develop an optimal compilation scheme.

To increase the performance of the generated code some deviations of the operational semantics of Clean [Eekelen 89] are made. Although Clean graphs are mapped directly to graphs in the graph store of the ABC -machine this store is not updated after every rewrite step. During the rewriting the information is stored on the stacks and (implicit) in the instruction sequence executed. The performance of ABC-programs heavily manipulating the graph store for simple computations will be low compared to ABC-code that employs the B -stack whenever appropriate. One of the ways to use the B-stack is outlined.

Section 5.1 contains an informal description of graph rewriting in Clean. In section 5.2 graph rewriting on the ABC-machine and the calling conventions are described informally. In section 5.3 the AST to hold the rewrite rules is defined. The code produced by the translation scheme is ABC-assembler which was introduced in chapter 3. The basic translation scheme is described in 5.4. In section 5.5 the runtime system is discussed. The way curried functions are implemented is outlined in section 5.6. Finally, a number of optimizations is discussed. The most important optimization is the use the B-stack to pass strict arguments of a basic type to functions.

### 5.1 Graph rewriting in Clean

To understand the code generation schemes discussed in this chapter it is important to have a proper view of the operational behaviour of a Clean program. A complete description of Clean is outside the scope of this work, we refer to [Barendregt 87b, Brus 87, Eekelen 89]. The semantics of Clean will be explained informally below.

Clean is a lazy higher order functional programming language based on Functional Graph Rewriting Systems, it is designed as intermediate language between arbitrary functional languages and (sequential) machines. A Functional Graph Rewriting System ( $F G R S$ ) is a Graph Rewriting System using the functional reduction strategy [Eekelen 88]. A reduction strategy is a function indicating the redex to be rewritten. The functional strategy delays the reduction of an expression until its value is needed. However, an expression is always reduced before its value is used. The functional strategy in Clean prescribes the same evaluation order as used in Miranda: lazy evaluation. A Graph Rewriting System (GRS) is an extension of Term rewriting Systems where the terms are replaced by directed graphs in order to avoid the duplication of work via sharing of expressions [Barendregt 87a, 87b, 88]. A Term Rewriting System (TRS) is a computational paradigm consisting of a collection of rewrite rules to transform terms (expressions) into equivalent terms [Klop 87].

A Clean program consists of a set of (typed) graph rewrite rules. The type system is based on the Milner/Mycroft type inference scheme [Milner 78, Mycroft 84].

A subgraph is a redex (reducable expression) if there is a left hand side (lhs or pattern) of rewrite rule that matches this graph. A match is a mapping from the pattern to the graph that is the identity on constants and preserves the node structure. A graph is in root normal form (rnf) if the whole graph is not a redex and will never become a redex. A graph is in normal form ( $n f$ ) if it does not contain any redex.

The rewrite rules are used to reduce the initial graph containing the symbol Start to normal form. The functional reduction strategy is used: rewrite alternatives are tried in textual order; patterns are matched from left to right; evaluation to root normal form is forced before an actual argument is compared with non-variable part of the pattern. An arguments is strict when its value always is needed in the reduction of the function. It is possible to deviate from the functional strategy by adding strictness annotations, denoted by I, to the rewrite rules. Such a graph is reduced eagerly to root normal form.

A redex is rewritten by constructing the graph specified in the right hand side ( $r h s$ ) of the rule: the contractum. Then all references to the root redex are redirected to the root of the contractum. There are also rewrite rule alternatives consisting of a redirection only; no contractum is specified in these rules. Nodes that cannot be reached from the root of the graph are garbage, they must be removed from the graph.

A small example is used to illustrate graph reduction in Clean. The example is even so small that no sharing of computations occurs. The Start rule initiates the computation of the length of a list. The rule Length takes two arguments; a number to record the number of list elements scanned and the list to be scanned.

```
:. Start -> INT ; II The type of the start rule
    Start -> Length O (Cons 3 (Cons 4 Nil)) ; Il The rewrite rule for Start
= Length ! INT I (List x) > INT ; Il Both argumen'ts are strict
    Length n (Cons a b) -> Length (+I n 1) b | |l +I is a delta rule to add
    Length n Nil }\quad->n={|2\mp@subsup{2}{}{\mathrm{ nd altemative;}}\mathrm{ a redirection
```

The reduction process is illustrated by the following rewriting sequence (the redex rewritten is underlined):
$\xrightarrow{\text { Start }}$ Length 0 (Cons 3 (Cons 4 Nil))
$\rightarrow$ Length ( +101 )(Cons 4 Nil)
$\rightarrow$ Length 1 (Cons 4 Nil)
$\rightarrow$
$\rightarrow \frac{\text { Length }(+111) \text { Nil }}{2}$

II a: This is the only redex, apply Start rule
II b: This graph as a whole is the new redex
$\| \mathrm{c}$ : The strict arguments are reduced first
$\| \mathrm{d}$ : Rewrite according to first altemative
|| e: Again one strict argument to be reduced
|| f : $1^{\text {st }}$ altemative does not match; use $2^{\text {nd }}$
$\| \mathrm{g}$ : This graph is in normal form
This rewriting process is depicted below. The graphs correspond to each of the steps shown above. The garbage that results from one rewrite step is drawn
gray, it is removed in the next snapshot. The dataroot is usually not shown, but it is included here to show the redirections clearly.


It is important to mention that Clean is a higher order language: symbols can be used Curried; without (some of) their arguments. The delta rule AP supplies one argument to a curried symbol, when all arguments of such a function are gathered its reduction is initiated. This is illustrated by the next example:

Twice $\mathrm{fx} \rightarrow \mathrm{AP} \mathrm{f}(\mathrm{AP} \mathrm{fx}) \quad ; \quad \|$ Twice applies the function f two times to the argument x Start $\quad \rightarrow$ Twice ++ I 0 ; $11++I$ is the delta rule to increment an integer

The reduction process is illustrated by the next rewrite sequence:

$$
\begin{aligned}
& \text { Start } \\
& \rightarrow \text { Twice ++I } 0 \\
& \rightarrow \quad \overline{A P++I(A P++I} 0) \\
& \rightarrow \quad+1(A P++10) \\
& \rightarrow \quad++\mathrm{I}(\underline{++10}) \\
& \rightarrow \quad+\quad+\text { I } 1
\end{aligned}
$$

II The initial graph
\| The graph after rewriting according to the start rule
II This graph will be rewritten by the delta rule AP
II AP has supplied the delta rule $++I$ its argument
II AP has supplied again the argument needed
II The delta rule ++ I must be applied again
\| The graph is reduce to its normal form

### 5.2 Graph rewriting on the ABC-machine

In this section an informal description of graph rewriting on the ABC-machine is given. The functional strategy and the rewrite algorithm are combined into a single piece of ABC-code for every rewrite rule. Associated with each rewrite rule there is a sequence of ABC -instructions which reduce a node in the graph to
its root normal form. Using the functional strategy a node must be reduced to mf , as soon as rewriting is initiated.
A graph reduction program in the ABC-machine consists of:

- rule dependent code to reduce a redex to its rnf according to the functional strategy and the rule in the system;
- code for the predefined rules;
- a run-time system: a fixed piece of code that initiates the reduction of the Start rule to normal form and prints the reduct.
For each Clean rule there is ABC-code to:
- prepare the arguments, i.e. construct a stack frame as expected by the code for the rule alternatives;
- match each rule alternative to the actual redex and rewrite accordingly;
- handle the situation that none of the rule alternatives is applicable. Code according to an additional rule alternative is generated for this purpose.
These pieces of code and their calling conventions (the interface between caller and callee) will be described below. Arguments and results can be passed among functions in the graph, on the A -stack and on the B -stack. It is more efficient to handle references to objects on the A-stack instead of in nodes of the graph: the references are directly available on the stack. When appropriate it is even more efficient to pass objects on the B-stack: the objects themselves are handled instead of references to the objects. The B-stack is not used in the basic compilation scheme, the description of its use is delayed until section 5.7.

To show the effect of the described actions a running example will be used. The reduction of a node containing the symbol $F$ and two matching arguments will be shown. The rewrite rules used are:


```
\(=\mathrm{G}(\) List x\()\) (List x\() \rightarrow\) INT ;
    Gab \(\quad \rightarrow 1\);
```

The overall assumption is that once initiated the reduction of a graph continues until it is in mf. Each rule alternative expects an A-stack frame containing a reference to the redex currently reduced and the node-id's of all arguments. The reference to the first argument is on top of the stack.

## Construction of the stack-frame

The reduction according to the rewrite rule can be initiated in two different ways. The preparation of the arguments needed depends on the calling circumstances.

The first situation is the initiation of the reduction by a jsr_eval instruction. In this situation the stack-frame consists only of a reference to the redex (see the definition of this instruction in chapter 3). To construct the desired stack frame
the node is marked as being under reduction, the arguments are pushed on the stack and the strict ones are reduced to rnf. The start of this code segment is indicated by the node entry.

In the second situation the reduction is initiated by a curried application of the function. Here, the node-id of the redex and the arguments will be on the stack, they are gathered by the code for the AP-rule. Strict arguments must still be reduced, so the code for the apply entry hooks up there in the code to prepare arguments.

In the figures below the relevant parts of the graph store and the A-stack frame are shown for a rewrite according to the rewrite rule above. New parts in the pictures are printed bold. Graphs with unknown contents are drawn as grey boxes.


Fig 5.3.a

b

c
a: The A-stack frame and graph upon entrance of the node entry.
b: The stack frame is extended with the node-id's of the arguments at the the apply entry.
c: The state expected by the rule altematives; both arguments are strict and hence reduced to mf .

## Rule alternative entries

For each alternative there is a separate rule alternative entry. The code for each entry consists of a matching phase which determines whether this alternative is applicable and a rewriting phase which performs the rewrite. When an alternative appears to be inapplicable, execution proceeds with the next one. Due to the generation of an additional alternative there is always a next one.

## Matching

Arguments are matched from left to right. If the formal argument is a variable the argument trivially matches. Otherwise, the argument is brought in root normal form and the symbol in the graph is compared with the specified symbol. When the symbol appears to match, its sub-arguments are pushed on the A-stack and matched themselves. If one of the arguments does not match the pattern, all sub-arguments are popped off the stack and execution proceeds with matching the next rule alternative. When the matching of all arguments succeeds the rewrite must be performed.


Fig 5.3.d

e

f
d: The situation after matching the top level of the first argument and pushing its subarguments.
e: The second sub-argument of the first argument is reduced and matched. After the successful match the node-id's of the sub-arguments are pushed on the stack.
f: The second argument also matches. So, this rule altemative is applicable.

## Rewriting

To increase the performance of the generated code some deviations of the operational semantics of Clean [Barendregt 87] are made. According to the semantics of Clean the contractum must be constructed (if present in the rule) and all references to the redex must be redirected to the contractum. Redirections are conceptually elegant and explain the semantics of graph rewriting very well. A straightforward, but also inefficient, implementation of redirection would have to examine the whole graph; every reference to the root of a redex is substituted by a reference to the contractum. A more efficient implementation is achieved by overwriting the redex with the root of the contractum. Although Clean graphs are mapped directly to graphs in the graph store of the ABC-machine this store is not updated after every rewrite step, but only when a root normal form is reached. During the rewriting the information is stored on the stacks and (implicit) in the instruction sequence executed. The redirection is implemented efficiently by overwriting the root node of the redex with the root of the contractum. In this way all references to the redex are automatically redirected to the contractum. Three situations are distinguished; the contractum is the next redex, the specified contractum is a root normal form and no contractum is specified in this rewrite alternative.

## Redirection to a reducible contractum

When the contractum is not known to be in rnf, the rewrite rule associated with the symbol in the root must be applied. It makes no sense to fill the redex with the root of the contractum. The new rewrite rule will unpack this node immediately since the functional strategy will indicate this redex as the one to rewrite. So, it is possible to delay the update of the graph store until a mf is reached. The code for this kind of rewrites constructs the stack frame for the first alternative of the new rewrite rule and proceeds immediately with that altemative.

g: The shared node, f , is constructed in the graph store.
h: The stack frame is ready for matching the first rule altemative of $G$.
i: The reduction according to rule G has been performed; a root normal form is reached.

## Redirection to a contractum in root normal form

When the contractum is a rnf, the root node of the redex is overwritten with the root of the contractum. The node containing the redex is indicated by the bottom of the A-stack frame. Rewriting to root normal form is completed. See figure 5.3.i for an example.

By overwriting the redex with its mf all references to the redex are automatically redirected to the contractum; a very cheap and commonly used implementation of redirection. On the ABC-machine it is always possible to overwnte a node with arbitrary new contents. For an other machine without variable sized nodes the redex is usually overwritten by an indirection node; a node indicating where this node actually can be found. A serious drawback of indirection nodes is the testing on these indirections required, even if a node is known to be in mf.

## Redirection to an existing graph

An existing graph cannot be rerooted at some other node without searching the whole graph or using indirection nodes. Since we want to avoid both, a copy of the root node is made. The redex and the top node of the graph are potentially both shared, so copying the top node to the redex immediately might result in the duplication work. To avoid the duplication of work the graph is reduced to rnf before the root node is copied. The graphs remain equal modulo unravelling which is sufficient for term graph rewriting and so sufficient for Clean. The obtained graph is different from the one obtained according to ordinary graph rewrite semantics.

To illustrate these graph manipulations an example is shown.

```
= Tll(Listx) -> (List x) ;
    Tl (Consht) ->t ; llaredirection
```

= F-> List INT ;
F $\rightarrow$ Cons 1 F ;


Fig 5.4.a

b

c
a: Starting the reduction of the graph TIF. State at the node entry. The arrows from outside indicate sharing.
b: After the reduction of the argument; state at the rewrite entry.
c: After copying the top node. Both nodes contain the proper mf.

## Example of the rewriting process

As example a rewrite rule is considered that computes the length of a list. The equivalent ABC-program generated by the specification presented in section 5.4 is listed (comments are supplied by hand). This is an important example since it illustrates the code generated by the basic compilation scheme.

```
: Length | INT I (List x ) \(\quad \rightarrow\) INT
    Length n (Cons ab) \(\rightarrow\) Length ( \(+\mathrm{I} \cap 1\) ) b
    Length \(n\) Nil \(\quad->n\)
```

The corresponding ABC-code:

| Descriptor |  |  |  |
| :---: | :---: | :---: | :---: |
| "Length" | "a_Length" 2 "Length" |  | , Il The generated descriptor |
|  | Label |  | II The node entry: preparing the arguments |
| "n_Length",Set_entry |  | "_cycle" 0 | , I\| Mark node to detect cyclic computations |
|  | Push_args | 022 | , II Push the arguments |
| Label |  |  | II The apply entry: reduce strict arguments |
| "a_Length", Jsr_eval |  |  | , Il Reduce first argument to mf |
|  | Push_a | 1 | , II Copy second argument to top of stack |
|  | Jsr_eval |  | Il Reduce it to mf |
|  | Pop_a | 1 | - Il Pop duplicated second argument |
|  | Label |  | II Entry for first rule alternative |
| "Length1", | Eq_desc_arity | "Cons" 21 | Il Match second argument |
|  | Jmp_false | "Length2" | II Continue with next altemative if match fails |
|  | Push_args | 122 | , Il Push sub-arguments; |
|  | Push_a | 1 | Il Rewrite according to alternative 1; Push b |
|  | Jsr_eval |  | $1 /$ Reduce it |
|  | Create |  | Il Node for result of +I n 1 |


|  | Create |  | II Node for 1 |
| :---: | :---: | :---: | :---: |
|  | Filli | 10 | Il Fill this node |
|  | Push_a | 5 | Il Push n. It is known to be a mf |
|  | Jsr | "+I1" | II Reduction of +1 $n 1$ |
|  | Update_a | 15 | Il Adapt stack frame for call of Length |
|  | Update_a | 04 | I |
|  | Pop_a | 4 | II Remove old (sub-)arguments |
|  | Jmp | "Length1" | II Continue with alternative 1 of Length |
|  | Label |  | 11 Entry for second rule alternative |
| "Length2", | Eq_desc_arity | "Nil" 01 | II Match argument |
|  | Jmp_false | "Length3" | Il Continue with next altemative if match fails |
|  | Fill_a | 02 | II Rewrite according to alternative 2;copy node |
|  | Pop_a | 2 | II Remove arguments from stack |
|  | Pin |  | II Rnf reached |
|  | Label |  | II Generated entry for additional alternative |
| "Length3", | Jmp | "type_error" | II This must be a type error! |

In the figures below several snapshots of the ABC-machine state are shown during the reduction of a graph according to the code for the rewrite rule Length.


Fig 5.5.a
b


d
a: Initial state at node entry.
b: Initial state at first rewrite altemative entry.
c: After successful match of first rule altemative.
d: Reduction according to the first rule alternative. First, the last strict argument is reduced. The state before calling +11 to reduce first strict argument is shown.


### 5.3 The AST for Clean

The Clean program to be transformed is represented by an AST as proposed in the previous chapter. The construction (lexical analysis and parsing) of the syntax tree is not treated here, this is done by standard compiler construction techniques [Aho 86]. The AST is assumed to be annotated with type and strictness annotations for the code generation. In Clean these annotations may be present in the program, but they can also be derived automatically by a strictness analyser [Nöcker 90] and by type inferencing [Milner 87, Mycroft 84]. When a symbol is used in a strict context (it is annotated as strict), all its strict arguments are annotated as Strict. The generation of these annotations is not treated here.

To facilitate the description the rewrite rules are assumed to be converted to a standard format. All shared nodes in a right hand side (rhs) are defined by explicit definitions on the textual outermost level of that rhs. They are ordered such that a node is defined in textual order before it is used as an argument in another node definition. This is always possible unless there is a cycle in these node definitions. Such nodes on a cycle are annotated as OnACycle.

To reduce the size of the description only a limited number of basic types is treated. In this chapter only the type INT is covered, in appendix B also booleans are handled. All other basic types can be treated in an analogue way.


| arg | $:=$ | Nodeld Term | annots node_id annots node | \| II An argument is either a node-id, Il or a node. |
| :---: | :---: | :---: | :---: | :---: |
| symbol | $:=$ | Symbol | symbolid | 1 II Type symbols |
|  |  | Function | symbolid | \| II Symbols with associated rule |
|  |  | Constructor | symbolid | \| II Symbols without associated rule |
|  |  | INTval | num | \| || Values of the basic type INT |
|  |  | BOOLval | bool | I\| Values of the basic type BOOL |
| node_id | = | [char] |  |  |
| symbolid |  | [char] |  |  |

A number of access functions with the obvious semantics is defined on this syntax tree. Their definitions can be found in appendix B.

Unfortunately, the Miranda type system does not allow to write just node as second alternative for arg. In the Miranda type scheme the constructor must identify the type of the expression uniquely. So, an additional constructor must be inserted in the AST when an other syntactical clause is included.

### 5.4 Compilation

In this section a compilation scheme is presented which transforms Clean rewrite rules into ABC-assembler. We assume that the AST represents correct Clean programs. The generated $A B C$-instruction sequence reflects the sequence of actions required for a rewrite specified above. Since the relevant information is stored in annotations, each rewrite rule can be transformed to ABC-code separately. Section 5.2 contains an example of the code produced by this scheme.

```
compiler :: clean -> abc_assembler
compiler = concat.map compile
```

For each rewrite rule a descriptor and code corresponding to the entries discussed above are generated.

```
compile :: rewrite_rule -> abc_assembler
compile (TypedRule typealts alts)
= fun_descriptor lhs ++
    prepare_args lhs ++
    al_entries alts 1 ++
    gen_type_error symbol_id (# args) (# alts)
    where lhs = get_lhs (hd typealis)
        NODE annots nodeid (Function symbol_id) args = hd lhs
compile (ConsRule typealts alts)
= fun_descriptor lhs ++
    prepare_args lhs ++
```

alt_entries alts $1++$
fill_with_rnf symbol_id (\# args) (\# alts)
where lhs = get_liss (hd typealts)
NODE annots nodeid (Function symbol_id) args = hd Ins
compile (UntypedRule alts)
= fun_descriptor lins +
prepare_args Ihs +
alt_entries alts $1++$
fill_with_rnf symbol_id (\# args) (\# alts)
where lhs = get_lhs (hd alts)
NODE annots nodeid (Function symbol_id) args = hd Ihs
compile (TypeRule typealts)
= descriptors typealts

## Descriptors

Descriptor generation for functions:

```
fun_descriptor :: graph -> abc_assembler
fun_descriptor (NODE annots nodeid (Function symbol_id) args:r)
= [ Descriptor
symbol_id
(apply_label symbol_id)
(# args)
symbol_id
```

For each constructor symbol defined in a type definition, a descriptor is generated. Function symbols defined in a type definition are partial functions, the corresponding descriptor is generated along with the code for that partial function. Constructors can also be used curried, so an apply entry containing fill code is defined for them.

```
descriptors :: rule -> abc_assembler
descriptors (alt:rest)
= cons_descriptor rhs ++
    fill_code rhs ++
    descriptors rest
    where rhs = get_rhs alt
descriptors []
= []
cons_descriptor :: graph -> abc_assembler
cons_descriptor (NODE annots nodeid (Constructor symbol_id) args:r)
= [ Descriptor
symbol_id
apply_lab
arity
```

```
                symbol_id
    where apply_lab
        = rnf_lab , if arity = 0
    = apply_label symbol_id , otherwise
arity = # args
cons_descriptor (NODE annots nodeid (Function symbol_id) args:r)
= []
fill_code :: graph -> abc_assembler
fill_code (NODE annots nodeid (Constructor symbol_id) args:r)
= [ Label (apply_label symbol_id)
    Fill symbol_id arity rnf_lab arity
    Rin
        ], if arity >0
= []
    , otherwise
    where arity = # args
fill_code (NODE annots nodeid (Function symbol_id) args:r)
= []
```


## Preparing the arguments

Two entries are generated to prepare arguments. The node entry is the regular entry; the target of the jsr_eval instruction. The apply entry is used by curried function applications.

```
prepare_args :: graph -> abc_assembler
prepare_args (NODE annots nodeid (Function symbol_id) args: alts)
= [ Label (node_label symbol_id)
    Set_entry cycle_lab 0
    pushargs 0 arity
    [ Label (apply_label symbol_id) ] ++
    reduce_strict_args args 0
    where arity = \# args
```

pushargs is a function which generates the appropriate instruction when there are arguments to push (arity $>0$ ).

Strict arguments are reduced by generating a Jsr_eval for each argument which is annotated Strict in the type definition. The node entry is the target of the jsr_eval instruction, it assumes that the root node-id is on top of the A-stack, so it must be copied if it is not at the top.

```
reduce_strict_args :: arg's -> a_src -> abc_assembler
reduce_strict_args (arg:args) n
= reduce_arg n ++ reduce_strict_args args (n+1) , if is_strict arg
= reduce_strict_args args (n+1) ,otherwise
reduce_strict_args [] n
= []
```

```
reduce_arg :: a_src -> abc_assembler
reduce_arg a_src
\(=[\) Jsr_eval \(]\), if a_src = 0
= [ Push_a a_src
    Jsr_eval
    Pop_a 1
] ,otherwise
```

The code for the first rule alternative is generated next to the code for the apply entry, no instructions are required to achieve the continuation.

## Rewrite alternative entries

The sequence of alternatives is decomposed to single alternatives. The alternative number is supplied to generate unique labels.

```
alt_entries :: rule -> num -> abc_assembler
alt_entries (alt:rest) alt_number
= alt_entry alt alt_number ++ alt_entries rest (alt_number+1)
alt_entries [ ] alt_number
= []
```

For each rule alternative, matching code and rewrite code will be generated. The code generated for the rhs assumes that the matching succeeded. If the graph does not match this lhs, execution proceeds with the next alternative and the rewrite code is not executed.
alt_entry :: rulealt -> num -> abc_assembler
alt_entry (Rewrite lhs rhs) alt_number
= match_code ++ contractum rhs bindings states asp
where (match_code,bindings,states,asp) = match Ihs alt_number
alt_entry (Redirect Ihs annots nodeid) alt_number
= match_code ++ redirection annots nodeid bindings states asp where (match_code,bindings,states,asp) = match Ihs alt_number

## Matching

Many of the subsequent transformation rules not only produce code, but also information of the contents of the stack frame. A compile-time number, called asp (for $A$-stack pointer), is used to record the size of the stack frame. The mapping from node-id's to offsets in the stack frame is recorded in the function bindings. The state of the referenced node is recorded in order to generate more efficient code.

| bindings | $==$ | node_id $\rightarrow$ offset |
| :--- | :--- | :--- |
| offset | $==$ | II Associates an index in the A-stack frame to a node-id |
| asp | $==$ | num |$\quad$ Il An index in the A-stack frame



The function match_args decomposes the matching into matching arguments one by one. A separate offset of the elements in the stack frame is passed since not every argument has a node-id, hence bindings cannot be used. For each specified pattern the evaluation of the actual argument is forced if necessary, match_arg takes care of the actual matching of the sub-graph.

```
match_args :: arg's -> bindings -> states \(->\) asp \(->\) label \(->\) offset \(->\) arity ->
                                    (abc_assembler,bindings,states,asp)
match_args (Nodeld an nodeid:rest) bindings states asp next offset arity
    = match_args rest bindings states asp next (offset-1) arity
match_args (Term ans (NODE n_ans id sym args):rest) binds states asp next oflset arity
= (code, bindings", states", asp")
    where code = bring_in_rnf annots a_src ++ match_code + match_rest
        a_src \(=\) asp-offset
        (match_rest,bindings",states",asp")
            = match_args rest bindings' states' asp' next (offset-1) arity
        (match_code,bindings',states',asp')
            = match_arg (NODE (ans ++ n_ans) id sym args)
                                    binds (record id InRnf states) asp next a_src arity
match_args [ ] bindings states asp next_alt offset arity
    = ([ ],bindings,states,asp)
```

The transformation rule match_arg generates code to check whether an actual argument matches the specified pattern. The actual argument is known to be in root normal form at this point. It is tested whether the symbol matches the symbol of the pattern. If this test fails, the rule does not match, all sub-nodes are
popped off the stack and the execution proceeds with the next rule alternative. If the pattern is not a basic value the sub-arguments are pushed and matched.

```
match_arg :: tree -> bindings -> states -> asp -> label -> a_src -> arity ->
(abc_assembler,bindings,states,asp)
match_arg (NODE anns id (INTval i) [ ]) binds states asp next a_src arity
= (code ,binds, states, asp)
    where code
                = typecheck ++
                [ EqI_a ia_src ] ++
                escape_false (asp-arity) next
                typecheck
            = [] ,if member anns IsINT
            = [ Eq_desc int_symbolid a_src ] ++
                        escape_false (asp-arity) next ,otherwise
match_arg (NODE annots id symbol args) binds states asp next a_src arity
= (code, bindings", states", asp")
where code
```

                    \(=[\) Eq_desc_arity (symbol_id symbol) n_arity a_src] ++
                        escape_false (asp-arity) next ++
                        pushargs a_src n_arity ++
                        match_sub_args
    bindings' = bind_args args asp' binds
states' $=$ record_args args states
asp' $=$ asp + n_arity
n_arity = \# args
(match_sub_args,bindings",states",asp")
= match_args args bindings' states' asp' next asp' arity
escape_false $n$ next generates code to pop the sub-arguments from the stack and a jump to the next alternative if the match failed.

```
escape_false :: nr_args -> label -> abc_assembler
escape_false nr_args label
= [ Jmp_false label ] , if nr_args =0
= [ Br_true 2
    Pop_a nr_args
    Jmp label ] , otherwise
```


## Rewriting according to a rule alternatlve with contractum

The contractum is built in a number of phases. First the shared nodes defined by the node-definitions are constructed. Then the root node is handled.

```
contractum graph -> bindings -> states > asp > abc_assembler
contractum (top node_defs) bindings states asp
= build_shared_nodes ++ root top bindings' states' asp'
    where (build_shared_nodes,bindings ,states',asp')
    = shared_nodes node_dels (bind top_ıd 0 bindings) states asp
    NODE annots top_id symbol args = top
```

The shared nodes laying on a cycle are created first. This enables the use of their node-1d's when the cycle nodes must be constructed Then the nodes are filled from top to bottom.

```
shared_nodes [tree] -> bindings -> states ->asp->(abc_assembler,bindings,states,asp)
shared_nodes node_defs bindings states asp
= (code, bindings", states", asp")
    where code = create_cycle_nodes ++ fill_nodes_code
        (create_cycle_nodes,bindings',states',asp')
                            = cycle_nodes node_defs bindings states asp
        (fill_nodes_code,bindings ',states",asp")
            = fill_nodes node_defs bindıngs' states' asp'
cycle_nodes [tree] -> bindings -> states -> asp -> (abc_assembler,bindings,states,asp)
cycle_nodes [] bindings states asp
    = ([ ],bindings,states,asp)
cycle_nodes (NODE annots nodeıd sym args rest) bindings states asp
= (code, bindings", states', asp ') , If member annots OnACycle
= cycle_nodes rest bindings states asp ,otherwise
    where code = [ Create ++ code_rest
        bindings' = bind nodeid asp' bindings
        states' = record nodeid Created states
        asp = asp +1
        (code_rest,bindıngs",states",asp") = cycle_nodes rest bindings' states' asp
```

The function fill_nodes takes care of the filling of the shared nodes in the order they are specified If necessary, a node is created

```
fill_nodes [tree] -> bindings -> slates -> asp -> (abc_assembler,bindings,states,asp)
fill_nodes [] bindings states asp
    = ([ ],bindings,states,asp)
fill_nodes (node rest) bindings states asp
= (code, bindings", states", asp")
    where code = creation ++ lill_code ++ fill_rest_code
        creation
    = [] , if previously_defined
    = [Create ] ,otherwise
        bindings'
    = bindings , if previously_defined
    = bind nodeid asp' bindings .otherwise
```

| $\quad$$=$ asp <br>  $=$ asp +1$\quad$, if previously_defined |  |
| :--- | :--- |
| $=$ asp' - bindings' nodeid | , otherwise |

The function fill_node fills an existing node with the specified value. It is used to fill shared nodes as well as arguments. Filling a node with a basic value is simple. Otherwise, code to construct the arguments is generated before the fill instruction. When a strictness annotation is present in the node, the function fill_strict is called to fill the node with a root normal form.

```
fill_node :: tree -> bindings -> states -> asp -> a_src -> (abc_assembler,states)
fill_node (NODE annots id (INTval i) [ ]) bindings states asp a_src
= (code, states')
    where code = [ Filli ia_src
        states' = record id InRnf states
fill_node (NODE ans nid (Function sid) args) binds state asp asrc
= fill_strict (NODE ans nid (Function sid) args) binds state asp asrc, member ans Strict
= (code, states") ,otherwise
    where code = build_args_code ++ [ Fill sid arity (node_label sid) (asrc+arity) ]
        (build_args_code,states") = build_args args binds states' asp
        states' = record nid Unknown state
        arity = # args
fill_node (NODE annots nid (Constructor sid) args) binds states asp asrc
    = (code, states")
    where code = build_args_code ++ { Fill sid arity rnf_lab (asrc + arity)
        (build_args_code,states") = build_args args binds states' asp
        states' = record nid InRnf states
        arity = # args
```

A node is filled with a root normal form by constructing the appropriated stack frame and calling the first alternative of the specified function. Using this entry, instead of the apply entry, results in a bit more code; the code for the evaluation of strict arguments is duplicated. However, the code produced is more efficient; when a strict argument is known to be in root normal form no attempt to reduce it , is necessary.

```
fill_strict :: tree -> bindings -> states -> asp -> a_src -> (abc_assembler,states)
fill_strict (NODE anns nodeid (Function sid) args) bindings states asp 0
= (code, new_states)
    where code = build_args_code ++ [ Jsr (reduction_label sid)
        new_states = record noderd InRnf states'
        (build_args_code,states') = build_args args bindings states asp
```

fill_strict (NODE annot id (Function sid) args) bindings states asp asre
$=$ (code, states)
where code

fill_strict node bindings states asp a_src
= fill_node node bindings states asp a_src

## Building arguments

Arguments are constructed by Build_args, the first argument specified is built last. So, it will be on the top of the A-stack as required.

```
build_args :: arg's -> bindings -> states -> asp ->(abc_assembler,states)
build_args [] bindings states asp
= ([ ],states)
build_args (arg:args) bindings states asp
= (code, states")
    where code = args_code ++ arg_code
        (args_code,states') = build_args args bindings states asp
        (arg_code,states") = build_arg arg bindings states' (asp+#args)
build_arg :: arg -> bindings -> states -> asp -> (abc_assembler,states)
build_arg (Nodeld annots nodeid) bindings states asp
    = (push ++ evaluation , new_states ) , if reduction_needed
= (push .states ),otherwise
        where new_states = record nodeid InRnf states
            push =| Push_a (asp - bindings nodeid)
            evaluation =| Jsr_eval
            reduction_needed = member annots Strict & ~(states nodeid=InAnf)
build_arg (Term annots node) binds states asp
= (creation ++ eval_code , strict_states ) , if member annots Strict
= (creation ++ fill_code , fill_states ),otherwise
    where creation = [ Create ]
        (eval_code,strict_states) = fill_strict node binds states (asp+1) 0
    (fill_code,fill_states) = fill_node node binds states (asp+1) 0
```


## Build contractum and overwrite root

Finally, the root of the contractum is constructed. If it is a constructor a root normal form is reached. After building the result a rtn instruction is generated to return to the initiator of the reduction. If the node contains a function symbol, the stack frame is adapted to the one required by the new function. A new stack frame for a function is created by building it on top of the current stack frame
and than copying it to the bottom. The new function is invoked by a jump instruction. In this way tail recursion is removed for free!

```
root :: tree -> bindings -> states -> asp -> abc_assembler
root (NODE annots id (INTval i) [ ]) bindings states asp
= pop_args asp ++
    [ Filli io
        Rtn
root (NODE annots nid (Constructor sid) args) bindings states asp
= build_args_code ++
    [ Fill sid arity rnf_lab (asp+arity) ] ++
    pop_args asp ++
    [ Rtn ]
    where (build_args_code,states') = build_args args bindings states asp
        arity = # args
root (NODE annots nid (Function sid) args) bindings states asp
= build_args_code ++
    clean_up (# args) asp ++
    [ Jmp (reduction_label sid) ]
    where (build_args_code,states') = build_args args bindings states asp
```


## Rewriting according to a redirection

A redirection is implemented by coping the root of the graph indicated to the redex. If it is known at compile time that this graph is in rnf it is just copied, otherwise it is reduced first in order to prevent the duplication of work.

```
redirection :: annots -> node_id -> bindings -> states -> asp -> abc_assembler
redirection annots nodeid bindings states asp
\(=1\) Fill_a newroot asp
    Pop_a asp
    Rtn ] , if states nodeid \(=\ln R n f\)
= update_astack newroot (asp-1) +
    pop_args (asp-1) ++
    [ Jsr_eval
        Fill_a 01
        Pop_a 1
        Rtn ] , otherwise
    where newroot = asp-(bindings nodeid)
```


## The additional rewrite alternative entry

A last entry is generated when the function has arguments; a rule alternative without arguments will always match, hence an additional entry is not needed.

```
fill_with_rnf :: symbolid -> arity -> num -> abc_assembler
fill_with_rnf symbol_id arity nr_alts
= []
= | Label (alt_label symbol_id (nr_alts+1))
    Fill symbol_id arity rnf_lab arity
    RT
```

gen_type_error :: symbolid -> arity -> num -> abc_assembler
gen_type_error symbol_id arity nr_alts
$=[]$
$=[$ Label (alt_label symbol_id (nr_alts +1 ))
Jmp type_error

### 5.5 The run-time system

The compilation scheme presented here produces code for each rewrite rule to reduce a graph to rnf. The run-time system contains the routine init_graph to build the start node and initiates its reduction to normal form. The reduction to normal form and the printing of sub-graphs in normal form is done by the routine _driver. Furthermore the descriptors for basic types and some general entries are defined. The _rnf entry is used for graphs in root normal form. The _cycle entry is stored in a node currently reduced. In this way it is detected that such a node is revisited, the functional strategy cannot find a rnf, hence an error message is generated. The type_error entry is used when none of the rule alternatives is applicable.

| [ | Descriptor |  |  |
| :---: | :---: | :---: | :---: |
|  | "_rnf" 0 "integer" |  | \|| Must be first descriptor!! |
|  | Descriptor |  |  |
| "BOOL" | "_rnf" 0 "boolean" |  | Il Must be second descriptor!! |
|  | Jmp | "init_graph" | \|| Must be first instruction!! |
|  | Label |  | \\| The initiator of graph reduction |
| "init_graph", | Create |  | II Create and fill the start node |
|  | Fill | "Start" 0 "n_Start" 0 . |  |
|  | Jsr | "_driver" | Il Print the $n f$ of start node |
|  | Print | "n" | I\| Print a newline |
|  | Halt |  | \|| Finished! |
| "_driver", | Label | 0 | \\| The global print driving routine |
|  | Pushl |  | II Closing bracket count |
|  | Label |  | II Label for tail recursion |
| "_print", | Jsr_eval |  | $1 /$ Reduce top node to mf |
|  | Get_node_arity 0 | 0 | $\\|$ Get number of arguments in node |
|  | EqI_b 00 |  | II No arguments? |
|  | Jmp_false | "_args" |  |


| "_print_last", | Label |  | 11 Print last argument |
| :---: | :---: | :---: | :---: |
|  | Print_symbol | 0 | , \\|| Do it |
|  | Pop_a | 1 | , Il Remove node |
|  | Pop_b | 1 | , Il Remove arity |
|  | Label |  | \|| Print the closing brackets |
| "_brackets", | EqI_b | 00 | - I\| Bracket count equal 0 ? |
|  | Jmp_true | "_exit" | , 1 Yes; finished |
|  | Print | $\left.{ }^{\prime \prime}\right)^{\prime \prime}$ | , Il No; print bracket |
|  | DecI |  | , I\| Decrement bracket count |
|  | Jmp | "_brackets" | , II Next bracket |
|  | Label |  | II Finished with this node. |
| "_exit", | Pop_b | 1 | , I\| Remove bracket count |
|  | Rtin |  |  |
|  | Label |  | II Start the printing of arguments |
| "_args", | Print | " ${ }^{\prime}$ | , II An opening bracket |
|  | Print_symbol | 0 | , \|| The Symbol of the node |
|  | Get_desc_arity | 0 | - Il Arity corresponding to symbol |
|  | Repl_args_b |  | , Il Replace node by its arguments |
|  | Pop_b | 1 | , II Pop descriptor arity |
|  | Label |  | Il Loop to print arguments |
| "_arg_loop", | Print | " ${ }^{\text {\% }}$ | , II Space between elements |
|  | EqI b | 10 | , Il Last argument ? |
|  | Jmp_false | "_next_arg" |  |
|  | Pop_b | 1 | , \|l Remove arg counter |
|  | IncI |  | , Il Increment bracket counter |
|  | Jmp | "_print" | , II Optimized tail recursion |
|  | Label |  | Il Print an argument; not the last one |
| "_next_arg", | Jsr | "_driver" | - Il Recursion to print argument |
|  | DecI |  | , Il Decrarg count; driver removes arg |
|  | Jmp | "_arg_loop" | , Il Next argument |
| "_mf", | Label |  | Il This graph is already in mf; |
|  | Rin |  | , \\|| Retum immediately |
|  | Label |  | \|| Functional strategy finds no mf! |
| "_cycle", | Print "Cyclic computationl Reduction interrupted", |  |  |
|  | Halt |  | - Il Stop the reduction |
|  | Label |  | II No rule alternatives is applicable! |
| "type_error", | Print "Type errorl Reduction interrupted", |  |  |
|  | Halt |  | II Stop the reduction |

### 5.6 Curried functions

The generic AP rule converts a curried symbol into its ordinary form as soon as all its arguments are available. Each curried symbol is a constructor since no rewrite rule can be applied, hence its context is _rnf. Each AP node provides one additional argument. When the number of arguments required by the symbol is collected, a stack frame according to the apply entry is constructed and execution
continues on that entry. Otherwise, the symbol and the arguments collected are stored in the node containing the AP.

| $\text { [ }{ }^{\text {AP" }}$ | Descriptor "a_AP" 2 "AP" Label |  | , |
| :---: | :---: | :---: | :---: |
| "n_AP", | Set_entry | "_cycle" 0 | - Il Mark node as being reduced |
|  | Push_args | 022 | , \|| Push arguments of AP |
|  | Label |  |  |
| "a_AP", | Jsr_eval |  | \|| Evaluate first argument; the function |
|  | Label |  |  |
| "AP1", | Get_node_arity | 0 | , Il Number of args in node |
|  | Get_desc_arity | 0 | , Il Number of args for cule |
|  | SubI |  | , Il Number of arguments needed |
|  | EqI_b | 10 | , Il More than 1 argument needed ? |
|  | Jmp_false | "args_needed" |  |
|  | Push_ap_entry | 0 | , I\| Push apply entry of function on C-stack |
|  | Get_desc_arity | 0 | - \|| Push arity |
|  | SubI |  | , Il Compute number of arguments in node |
|  | Push_b | 0 |  |
|  | Repl_args_b |  | - $\mid$ Construct top of stack frame |
|  | Pop_b | 2 | , \|| Clean up B-stack |
|  | Rin |  | - Il to apply entry pushed above |
|  | Label |  |  |
| "args_ne | ded" |  |  |
|  | Push_a | 1 | II Reference to curried function |
|  | Add_args | 113 | - II Add argument and overwrite AP node |
|  | Pop_a | 2 | - \\|l Clean up stacks |
|  | Pop_b | 1 |  |
|  | Pits |  | ] \|| Rnf reached |

This is illustrated by the following example:

```
Start -> AP Inc 2 ;
Inc -> AP Plus 1 ;
Plus -> +I
```

Several states of the reduction are depicted below.

a: Initial state. The reduction starts at the node entry of Start.
b: Creation of the arguments of AP. Second argument is built. First argument is strict.
c: Reduction according to Inc. Second argument is created, first argument is strict.
d: Reduction according to Plus. The redex node is filled with a mf.
e: +I has not enough arguments to become a function. So, the node is filled with the mf.
f: After reduction according to +I the top node is replaced by its mf.

### 5.7 Optimizations

The code generated by the compilation scheme presented in section 5.4 , can be improved at many points. The most important improvement is to use the B -stack instead of nodes to pass basic values between functions. The use of the B-stack is described informally and illustrated with an example. Afterwards some other optimizations are mentioned.

## The extended compilation scheme using the B-stack

Basic values are manipulated always on the B-stack in the ABC-machine. Every computation involving basic values requires the transportation of the values to the B-stack and the shipment of the result back to a node in the graph store. When the result is used again as argument in another computation there is much data transportation. To reduce the unnecessary movement of data, the compilation scheme must be changed such that basic values stay on the B -stack as much as possible.

To achieve this, the calling conventions for the rewrite altematives are changed: strict arguments of a basic type are passed on the B-stack and the result of a reduction is left on the B-stack when it is of a basic type. The calling conventions for the node entry and the apply entry remain unchanged. The code corresponding to the apply entry takes care of the transport of basic values between the graph and the B-stack for this function. After the reduction of a strict argument of a basic type it is transported to the B-stack. When the result of the function is of a basic type the first rule alternative entry is called as a subroutine. The mff produced on the B stack is transported to the node containing the redex. The complexity of the code generation is increased significantly by this new calling convention. Arguments and results must be moved to the desired place at every occurrence. This is not difficult, but involves an elaborated case analysis.

## An example of the use of the B-stack

To show the changes resulting from the new calling conventions the ABC-code for the Length example of section 5.2 is shown.

```
:: Length!INT!(List x) -> INT
    Length n (Cons ab) -> Length (+I n 1)b
    Lengthn Nil >
```

Both arguments are strict. The first argument is also a basic value. So, it will be passed on the B-stack to the rule alternatives. The result will also be passed on the B-stack by the rule alternatives. The rewrite entry for +1 also expects its arguments and leaves its result on the B-stack.

|  | Descriptor |  |  |
| :---: | :---: | :---: | :---: |
| "Length" | "a_Length" 2 "Length" |  | , II The generated descriptor |
|  | Label |  | Il The node entry: preparing the elements |
| "n_Length", | ,Set_entry | "_cycle" 0 | , Il Mark node to detect cyclic computations |
|  | Push_args | 022 | , Il Push the arguments |
|  | Label |  | II The apply entry: reduce strict arguments |
| "a_Length", | ,Jsr_eval |  | , Il Reduce first argument to mf. |
|  | Pushi_a | 0 | , Il Copy first argument to the B-stack. |
|  | Pop_a | 1 | , Il Pop first argument from A-stack. |
|  | Jsr_eval |  | , Il Reduce second argument to mf. |
|  | Pop_a | 1 | , Il Pop duplicated second argument. |
|  | Jsr | "Length1" | , Il Initiate rewriting by altematives. |
|  | Fill_b | 00 | Il Fill node with result of the reduction. |
|  | Pop_b | 1 | , ll Pop result of the reduction. |
|  | Rtn |  | , Il Done; node is in mf. |
|  | Label |  | II Entry for first rule alternative. |
| "Length 1 ", | Eq_desc_arity | "Cons" 20 | , Il Match arg 2. |
|  | Jmp_false | "Length2" | II Goto next altemative if match fails. |
|  | Push_args | 022 | , I\| Push sub-arguments. |
|  | Push_a | 1 | I\| Rewrite according to alternative 1; Push b |
|  | Jsr_eval |  | , Il Reduce it. |
|  | Pushi | 1 | , II Second argument for +l . |
|  | Push_b | 1 | , Il First argument for +1 . |
|  | Jsr | "+I1" | , \|l Reduction of +I $\cap 1$. |
|  | Update_a | 03 | , II Adapt A-stack frame for call of Length. |
|  | Update_b | 01 | , II Adapt B-stack frame for call of Length. |
|  | Pop_a | 3 | , Il Remove old arguments from A-stack. |
|  | Pop_b | 1 | , II Remove old arguments from B-stack. |
|  | Jmp | "Length1" | , HContinue with alternative 1 of Length. |
|  | Label |  | II Entry for second rule alternative. |
| "Length2", | Eq_desc_arity | "Nil" 00 | , II Match argument. |
|  | Jmp_false | "Length3" | , II Continue with altemative 3 if match fails |


|  | Pop_a | 1 | Il Remove arg. |
| :---: | :---: | :---: | :---: |
|  | Rta |  | 11 Rnf reached! |
|  | Label |  | 11 Generated entry for additional alternative. |
| "Length3", |  | "type_error" | 11 This must be a type error! |

In figure 5.7, snapshots of the reduction process corresponding to the pictures in Fig 5.5. are shown.

a: Initial state at node entry.
b: State at first rewrite altemative entry.
c: After successful match of first rule altemative.
d: Reduction according to the first rule altemative. State before calling +I1.

e: State retumed by +11.
f: State before calling Length1 recursively.
g: State before last call of Length1.
h: State of ABC-machine after the mf is reached.
This example shows the successful use of the B-stack, very few movement of data between the graph and B-stack is necessary. The next example shows that one is not always that lucky.

| $=11 x$ | $\rightarrow x$ |
| :---: | :---: |
| IX | $\rightarrow x$ |
| = F!INT! ${ }^{\text {NT }}$ | $\rightarrow$ INT |
| Fab | $\rightarrow+\mathrm{I}(\mathrm{l}$ a) b |

Both arguments and the result of function F will be passed on the B -stack by the rule alternative entry. The identity function expects its argument and leaves its result on the A-stack. So, a has to be put in a node to pass it to I and the result must be extracted of a node to pass it to +I .


Fig 5.8.a

b

c

d

e

f
a: Machine state at node entry of rule $F$.
b: State at rewrite altemative entry for $F$.
c: Second argument of +1 copied to the B -stack. Initiating the reduction of first strict argument.
d: The mf returned by I .
e: The stack frame to call + I1.
f: Finally, the mf is copied to the redex by the apply entry of $F$.

## Other optimizations

The code generated by the schemes above is more efficient than the scheme in 5.4 , but can be improved further. Here we mention some other optimizations possible.

Improve the performance of predefined functions (delta rules) by handcoding them in ABC-assembler.

The use of in-line code substitution for very short functions. For instance, the AddI instruction instead of a Jsr +11 .

Use information of predefined functions to improve the code generated by the compilation schemes. For instance, the predefined conditional function:

$$
\begin{aligned}
&= \text { IF I BOOL type type } \rightarrow \text { type ; } \\
& \text { IF True then else } \rightarrow \text { then } \\
& \text { IF False then else } \rightarrow \text { else ; }
\end{aligned}
$$

Obviously, IF is strict in its first argument. Its second and third argument are never both needed. In a strict context building the then and else part can be delayed until the condition is evaluated. Then it is clear which one is needed, only that one has to be built.

In the present compilation schemes the matching of each rule alternative starts from scratch. However, the use of information from the previous alternative and type information can speed up the matching process considerably. So, it is better to use a finite state machine for the matching of the rule alternatives instead of starting all over for each alternative.

Constant graphs which occur in the rhs of rewrite rules need not to be built each time the rule is used. These graphs can be created once. The node-id's of the root of these graphs can be used instead of building a new one at each occurrence.

Use a scratch node for nodes which will be used once for a very short time. This occurs when a function delivers its result in the graph which is required on the B-stack; the result is stored in the node, moved to the B-stack and the node becomes garbage immediately.

## Efficiency gain of the optimizations

To analyse the effect of the mentioned optimizations a measurement of the performance of the well known nfib function is presented. The integer delivered by the nfib function is the number of function calls done. The nfib-number is obtained by dividing this number by the execution time. So, the nfib-number is the number of function calls per second.


The table below gives an idea about the gain in efficiency obtained by the optimizations. However this example is not representative for all programs. The nfib number is extremely sensitive for the optimizations discussed.

These measurements where done on a SUN $3 / 280$ under SunOS 4.0.3, using the ABC-machine implementation made by Gé Weijers [Weijers 90]. Measurements are accurate within $10 \%$.

| Nfib-number |  |
| :---: | :--- |
| 27,000 | Conditions <br> Code generation according to the scheme presented here. No <br> strictness information nor type information is provided. At run- <br> time a type check on the argument of Nfib is performed. |
| 35,000 | Using type and strictness information for the delta rules only. In- <br> line code substitution is used for these delta rules. In the rest of the <br> bench-marks these delta rules are used. |
| 43,000 | Using type information at compile time; the type check on the <br> argument of the nfib function is dropped. <br> Using only the strictness information of the nfib function. No type <br> information is supplied. So, the type of the argument will be <br> checked at run-time. <br> Type and strictness information are supplied to the basic transla- <br> tion scheme. |
| 70,000 | Using type and strictness information in the extended translation <br> scheme. The integers are passed on the B-stack but no in-line code <br> is used for the delta rules. |
| 415,000 | Using type and strictness information in the extended translation <br> scheme. The integers are passed on the B-stack and in-line code is <br> used for the delta rules. |

495,000 Result of the same ABC-code, but using a better register allocation in the ABC -compiler [Groningen 90]. The ABC-code resulting from an nfib-function definition with a conditional is more efficient. The corresponding Nfib-number on the Sun is 571,000 on a Mac IIfx this nfib-number is $1,009,000$ !
2,000 Same code interpreted by the PABC simulator [Nöcker 89].
300,000 A similar nfib function written in C [Kernighan 78]. The nfib function written in Miranda; executed by the interpreter.
1 ABC-specification executed by the Miranda interpreter
10 ABC-specification without superfluous checks in the microinstructions, executed by the Miranda interpreter.

The code for Nfib in ABC-assembly using the extended compilation scheme becomes:

|  | Descriptor |  |  |
| :---: | :---: | :---: | :---: |
| "Nfib" | "a_Nfib" 1 "Nfib" |  | , II The generated descriptor |
|  | Label |  | Il The node entry. prepare arguments |
| "n_Nfib", | Set-entry | "_cycle" 0 | , \|| Mark node to detect cyclic computations |
|  | Push_args | 011 | , II Push argument on A-stack |
|  | Label |  | 11 The apply entry. |
| "a_Nfib", | Jsr_eval |  | , Il Reduce strict argument |
|  | Pushi_a | 0 | , Il Copy argument to B-stack |
|  | Pop_a | 1 | , Il Remove argument from A-stack |
|  | Jsr | "Nfib1" | , III Initiate calculation on B-stack |
|  | Filli_b | 00 | , I\| Copy result from B-stack to node |
|  | Pop_b | 1 | , I\| Remove result from B-stack |
|  | Rtin |  | II Rnf reached |
|  | Label |  |  |
| "Nfib1", | EqI b | 00 | , Il match rule 1 ; is argument 0 ? |
|  | Jmp_false | "Nfib2" | , II No; continue with second rule altemative |
|  | Pop_b | 1 | , Il Yes; remove arg |
|  | Pushi | 1 | , II Store result |
|  | Rm |  | , Il Done |
|  | Label |  |  |
| "Nfib2", | EqI_b | 10 | , Il Match rule 2; is argument equal to 1? |
|  | Jmp_false | "Nfib3" | , Il No; continue with rule altemative 3 |
|  | Pop_b | 1 | - Il Yes; remove argument |
|  | Pushi | 1 | , II Push result |
|  | Rton |  | - Il Done |


| "Nfib3", | Label |  | II Rule altemative 3; no matching required |
| :---: | :---: | :---: | :---: |
|  | Pushi | 2 | , Il Compute -I $n 2$ |
|  | Push_b | 1 | II Push n |
|  | SubI |  | II In-line code substitution for -I |
|  | Jsr | "Nfib1" | , II Compute Nfib (-I п 2) |
|  | Push_b | 1 | , Il Compute --I n; push argument |
|  | DecI |  | II In-line code substitution for --I |
|  | Jsr | "Nfib1" | - Il Compute Nfib (-I n) |
|  | AddI |  | II In-line code substitution for +I |
|  | Update_b | 01 | , Il Replace argument by result |
|  | Pop_b | 1 | - Il Remove copied result |
|  | IncI |  | II In-line code substitution for ++I |
|  | Pton |  | \| Il Done |

### 5.8 Discussion

This chapter presents an abstract compilation scheme to translate Clean to ABCassembler. This is an example of a non-trivial translation containing context sensitive decisions. Due to the nature of the translation to be performed the specification is rather elaborate, but it remains very readable.

It appears to be valuable that the specification is executable; the generated code can be observed and executed. Execution serves as a test for the specification, due to its nature it is very hard to verify its correctness. Using this prototype it is easy to glance at the code generated according to slightly different schemes; it appears to be easy to make small changes in the compilation schemes. Combining this prototype with the prototype of the ABC -machine presented in chapter 3 it is possible to execute the Clean programs. Important observations of the dynamically behaviour can and have be obtained in this way. Furthermore there are of course the usual advantages: the consistency can be partially checked by the Miranda compiler and the semantics of the description language are clear.

The specification is used as guide-line for the construction of the corresponding product [Smetsers 89]. This compiler uses the B-stack in the way described here and incorporates the first three other optimizations mentioned above as well as many other optimizations. Although it is in principle possible to describe each and every optimization in a specification this is not done. The specification obtained would be very large and, hence, almost as difficult to read as the actual implementation. The specification presented here shows clearly the kind of code generated and enables the informal description of many optimizations.

# Chapter 6 Artificial Neural Networks 

### 6.1 Description tools for artificial neural networks

6.2 The Adaptive Linear Combiner
6.3 The Adaline
6.4 The Perceptron
6.5 Learning Vector Quantization
6.6 Feedback Networks
6.7 Discussion

The popularity of artificial neural networks is rapidly increasing. These networks appear to be able to solve relatively easy some problems that are hard to solve with conventional algorithms. Especially pattern recognition tasks are performed very well by some networks. In order to detect which aspects of neural networks are responsible for this success, it is necessary to have a good understanding of the paradigm used in neural network computations. We will show that functional languages can be used to obtain a uniform description of various artificial neural networks. Such a uniform high-level description will make it easier to spot the essential differences and similarities among the neural network models used.

There are many possible descriptions of a biological neural network. Each description shows some aspects of the network. Some of the possible descriptions are: morphological, chemical (the description of the chemical reactions and matters involved, e.g. neurotransmitters), biological (description of the neurons as biological cells), physical (spike generation and travelling over the axons, or system description of large ensembles of neurons) and psychological (description of complicated behaviour). The missing item in this list is a description of the type of computations performed and algorithms used by a neural network. The nature of these computations is largely unknown.

One of the ways to explore the computational models used in the brain is to make a study of artificial neural networks. These artificial networks are models
inspired by the brain. One usually starts with very simple models that can perform only a very limited set of tasks. Other features, increasing the complexity, are introduced step by step when the current model cannot execute the next task. More than 50 networks are proposed in the literature [Hecht-Nielsen 87b, 88].

In this chapter a specification of some important artificial neural networks in a functional programming language is given. First, a set of description tools is defined. Then, each specific network is described with these primitives. For the description of artificial networks it is very convenient that the descriptions are executable. Due to the used nonlinear functions and stochastic terms a complete analytical study of the properties of these systems is very hard or impossible.

These descriptions show that the processing elements used in these networks are very similar. The use of these elements distinguishes the various networks. To indicate the huge simplifications made in artificial networks with respect to real networks a rough sketch of biological neural networks is given.

## Biological neural networks

The human brain consists of a very huge number (estimations range from 1010 to $1^{13}$ ) of neurons (brain cells), each one receives input from other neurons (typical about $10^{4}$ ) [Schadé 84]. Based on this input a short pulse on the output, called a spike, may be generated. This spike travels over the output axon and serves as input for other neurons. A spike is a pulse-shaped change of the potential over the axon membrane.

Nerve conduction is studied since the late 1940's. The shape and conduction speed along the axon of a spike are related to the observed changes in membrane permeability by the Hodgkin Huxley equations [Hodgkin 52]. The model does not explain why the membrane permeability for various ions changes; it describes just the consequences. These equations can be found in many textbooks, e.g. [Hobbie 78].

When a spike reaches the connection of the output axon with some input axon it changes the state of the corresponding neuron. In most situations chemical substances, called neurotransmitters, are released in the synaptic junction in quantized packets. These neurotransmitters change the interior potential of the post synaptic neuron, it becomes either more positive or more negative depending on the neuro transmitter released. If the potential becomes high enough, above the threshold value, a spike is generated on its output axon. Properties of neurotransmitters are studied since long, e.g. [Koopman 60].

There is morphological evidence that the number of synapses in the brain increases rapidly with the age. It is believed that this is one of the mechanisms to let the brain learn.

### 6.1 Description tools for artificial neural networks

Artificial neural networks are a crude simplification of biological networks. These models are most suited for relatively simple tasks, like sensory and motor functions. Each artificial neuron is a simple processing element which receives input from a fixed set of neurons. The output value is not continuously computed in the massive parallel way of the brain, but in discrete time steps. Artificial neurons do not communicate by spikes, but by a sequence of values representing neuron activity. A single scalar number represents the strength and state of a synapse. The input value is multiplied by this weight to obtain the contribution of this input to the potential. A learning procedure adjusts these weights to let the network perform the desired task. The number of processing elements used ranges from one to several thousands.

This chapter presents a set of description tools for artificial neural networks and uses them to specify some well-known networks. This approach does not always yield the most direct and compact description of a given network, but makes the similarities and distinctions between the various networks better visible. Here it is assumed that the network has a layered structure of uniform processing elements. Since the networks described here can be cascaded also networks with layers of different nature can be captured by this description method. The tools developed here are also used to describe non-layered networks [Koopman 90, Rutten 90].

Artificial neural networks are characterised by

- the values processed;
- the processing elements used;
- the topology of the network (interconnection of processing elements);
- the synchronous or asynchronous character of the computations; and
- the way the processing elements are tuned to let the network perform some task. Elements are either completely predetermined or some parameters are tuned in a learning phase.
First, a set of description tools is developed in a brief discussion of these characteristics. Then, some neural networks are described using these tools.


## The values processed

The values on the axons are numbers representing neuron activity. There are two types of domains used; either binary or continuous.

A binary activity quantity has one of two possible values; these values represent firing at maximum rate and quiet. The numerical values are usually 0 and 1 , or -1 and 1 . The latter variant has some advantages since most synapses compute a weighted sum of the input values, but every weighted sum of zeros yields 0 . Moreover, in many applications it is convenient that on and off of a neuron are in some sense symmetrical.

A continuous activity quantity can have any value within some limits as 0 and 1 , or -1 and 1 . These bounds have the same meaning as for binary values, but also intermediate states of neuron activity can be modelled.

Our description uses the type value for neuron activity, ordinary numbers are used to represent it. When appropriate, values are indicated as input, output or target value.

```
value == num
input_value =m value
output_value == value
target_value == value
```

A collection of simultaneously occurring values is modelled by a list of values called state.

```
state == [value]
input_state == state
output_state == state
target_state = = state
```


## Neuron characteristics

In the simplest model a neuron takes an input state and yields an output value determined by this state. In more complicated models there may be additional parameters such as a noise input or a history term representing some dynamic properties. Most of the neural network models found in the literature do not need these extensions. Therefore, we will model our neurons by:
neuron $\quad==$ input_state $->$ output_value
Each neuron consists of a synapse which yields an internal potential, and a generator which produces the output value based on this potential. This element is introduced by McCulloch and Pitts [McCulloch 43] and has become well-known through the work of Hopfield [Hopfield 82]. Artificial neural networks usually employ this type of elements. This model forms an abstraction of biological neurons.


Fig 6.1 The McCulloch-Pitts model of a neuron.
The internal potential of an element is fully determined by the current input.

## The synapse

A synapse is a device (function) that computes the internal potential based on the input state. Usually, a weighted sum of the input values is used. The only difference between various neurons is the set of weights used to compute the potential. Learning is modelled by changing these weights. For these reasons the weights are passed as arguments to the synaptic function and to the elements.

```
weights == [num]
potential == num
synapse == weights }>>\mathrm{ input_state }->\mathrm{ > potential
```

Some well-known synaptic functions are:

```
wsum weights inputs = sum [weightsli * inputsli | i <- [0..#weights -1]] }\mp@subsup{}{}{1
neocognitron_synapse weights input
= (1+excitation)/ (1+inhibition) - 1
    where excitation = sum [weights \i* input \i|i<-[0..#weights-1]; weights \i>0]
        inhibition =- sum [ weights I |* input I i| i<- [0..#weights-1]; weights ! i<0]
```

When one views the sequence of weights and the input state as vectors, wsum just computes the inner product (scalar product) of these vectors.

## The generator

The generator function yields the output value determined by the potential.

```
generator = = potential -> output_value
```

[^6]When an index is used at several places in an expression we generate the sequence of indices rather than the sequence of elements chosen. Using the generation of list elements instead of the generation of indices, the definition above would become:

> sum [weight * input | (weight,input) <- zip2 weights inputs]

This generator is generally a nonlinear function of the potential. Some models involve a stochastic term in the generation of the output value. This is adequately modelled by adding a noise term to the potential. The most often used generator functions are a hard limiter for a binary output and a sigmoid or pseudo linear function for continuous output values. Some well-known generator functions are:

```
hard_limit v |l binary output
    = high, if v > threshold
    = low ,otherwise
threshold = (high + low) / 2
```

The optimum threshold value depends on the network model used and the specific problem at hand.

```
sgn v Il binary output
= 1, if v>0
=-1, otherwise
```

clip v
II continuous output, pseudo linear
$=$ low, if $v<m$ low
$=v$, if low $<v<=h i g h$
$=$ high , if high $<v$
sigmoid $v=($ high-low $) /\left(1+\exp \left(-s^{*} v\right)\right)+$ low $\|$ continuous output
s is a small positive constant determining the 'sharpness' of the function. For very large values of the sharpness, the sigmoid function resembles the step function.

The bounds high and low are separately defined for each specific network.
In many models the input of the generator is the difference of the potential and a threshold value, instead of the potential itself. Whenever the potential is determined by a weighted sum, which is by far the most used synaptic function, a connection to a fixed valued input can replace the threshold. For a generator function, $f$, the output is determined by

```
output = f(potential - threshold)
potential = wsum weights input_state
```

This is equivalent to

```
output = f potential'
potential' = wsum (-threshold: weights) (1: input_state)
```

The latter formulation is more convenient since no individual threshold values for the elements have to be maintained. Another benefit of this formulation is
that the threshold can be optimized by the learning algorithm used for the weights.

## Processing elements

The proper combination of synaptic function and generator function specifies the processing elements. Some well-known processing elements are:


The usual choice for the bounds high and low is 1 and -1 respectively. With these values the functions sgn and hard_limit become equivalent. This makes the processing elements used in the various networks even more similar.

## Neuron interconnection patterns and evaluation orders

Artificial neurons are usually arranged within layers. The layers consist of identical elements, the only difference between the various neurons is the vector weights used in the synapse. We prefer to specify a layer of neurons by the element function for all neurons and the weight vector for each element, instead of a sequence of neuron functions. It is much easier to adapt these weight vectors by a learning algorithm and to monitor their evolution than to change and monitor the evolution of a function.


In this description the size of the network is deduced from the sizes of the weights set determining the synapses in the network. Usually, the size of a network is not changed during its use, albeit it is possible that the leaming algorithms change the size of the network dynamically. The initial network is generated by the function initial_net which is defined for each network separately. The size of the generated net is determined by the list of numbers.

```
initial_net :: [num] -> weights_net
```

To avoid computational problems resulting from cyclic dependencies, it is usually assumed that there are only unidirectional connections between two adjacent layers. The units in the first layer, called the input layer, receive the input state
and yield a state to be processed by the next layer etcetera. Finally, the last layer, the output layer, produces the response of the network.

The output state of a layer of neurons is specified by a function parametrized by the element used, and the list of weight vectors for each individual element. This function takes the input state as argument and yields the corresponding output state.

> eval_layer $::$ element $\rightarrow$ weights_layer $->$ input_state $->$ output_state
> eval_layer element layer input = [ element weights input $\mid$ weights <- layer]

With a single layer of neurons only very simple networks can be built. Using a cascade of these layers, a network with better properties can be built: a multilayer feed forward network. A list of representations of layers represents a layered network. Here it is assumed that all layers contain identical elements. A network consisting of layers with different elements in the layers can be described by a cascade of the networks described here. The constant neurons, used to represent the threshold, are added to each layer. The number of constant neurons is determined in the specification of a specific network.

```
feed_forward :: element -> weights_net -> input_state \(->\) net_state
feed_forward element weights input
= state
        where state
            \(=\) (input ++ constant_neurons): Il state!0 is the input state
                    [eval_layer element (weightsII) (statell) ++ constant_neurons
                            | \(1<-\) [0..\#weights -1]]
constant_neurons :: state
constant_neurons = rep nr_constant_neurons high
select_output :: net_state -> output_state
select_output state
\(=\) output_state \(\quad\), if nr_constant_neurons \(=0\)
= take (\#output_state - nr_constant_neurons) output_state, otherwise
    where output_state = last state
```

The input state is recorded in the state of the network since it is needed in several learning algorithms. Unfortunately, this additional vector in the state of the network has as consequence that the element determined by the vector at weightsIIII corresponds to statel( $(1+1)$ !i.

Some artificial neural networks do have connections within one layer, the introduction of a delay solves the computational problems due to the cyclic dependencies. The neuron uses the output produced some time steps ago as input, instead of the current output. In these feedback networks an input state is presented to the network and the output is used as new input state. This requires that the input state and output state have the same size. These networks consist
usually of a single layer of elements. For feedback networks it matters whether the elements are updated synchronously or asynchronously. When the elements are evaluated asynchronously, i.e. one by one, the element that is updated as last one, 'sees' almost the new state. When the elements are synchronously evaluated all elements have the same input state. The function feed_back_syn searches a stable state in a sequence of states generated by an iteration of feed_forward passes over the initial state.

```
feed_back_syn :: element -> weights_net -> input_state -> [output_state]
feed_back_syn element weights input
= to_limit (iterate (select_output.feed_forward element weights) input)
```

The funcion to_limit takes the leading part of a list. The list is broken when two consecutive elements are equal. It is defined as:

```
to_limit :: [*] -> [']
to_limit (a:b:x)
= [a] , if a=b
= a : to_limt (b:x) , otherwise
```

The function feed_back_asyn implements an asynchronously updated single layer network, the elements are updated one by one in a pseudo random order.

```
feed_back_asyn:: element -> weights_net -> input_state -> [output_state]
feed_back_asyn element [weights] input
= to_limit (iterate (eval_one_by_one element weights) (input ++ constant_neurons))
eval_one_by_one :: element -> weights_layer -> input_state -> output_state
eval_one_by_one element w_layer input
= eval_neurons element w_layer input [0..# w_layer -1]
eval_neurons :: element -> weights_layer -> state -> [num] -> state
eval_neurons state element w_layer [] = state ++ constant_neurons
eval_neurons state element w_layer indices
= eval_neurons new_state element w_layer rest
    where n = random (1993+entier (sum state)) mod # indices Il pseudo random
        rest = remove n indices |l remove this index from the list of neurons
        index = indices In |l index of neuron to update in this iteration
        new_state = update index new_value state
        new_value = element (w_layerlindex) state
```

There are several successful networks using a topology not covered by the class of layered networks described above, e.g. the Boltzmann machine is a randomly connected network. A sparsely connected network can sometimes be modelled adequately by a fully connected network where some of the weights are kept zero.

The 'knowledge' of a neural network is stored in its architecture and the weights used in the synapses. Generally, the desired weights are unknown. The function to be performed is only specified by a set of input-output pairs. The weights cannot be determined immediately, but must be deduced in some way from the samples. There are several algorithms known to obtain a reasonable set of weights by training the network. These learning rules can be classified into

- Unsupervised learning. These learning rules must set the weights based upon a set of sample inputs. No supervisor system is available to indicate whether the network behaves correct. This is used mainly in pattern recognition, the samples presented are the samples to be recognised later. So, sample and target are identical. The best-known unsupervised learning rule are the Hebbian-type learning ruls: the strength of a connection between two units is proportional to the correlation in activity. Activity in one unit is associated with the activity in the other unit. In this way the network learns to associate patterns of neural activity.
- Supervised learning. In these situations there is a teacher which presents sample patterns and the desired output state. The general principle is to feed a network a sample input and target. The weights in the network are adapted based on the difference between target and output state, such that the response of the network is improved when the same sample is presented to the network again. Using the target it can be detected which output units are wrong. For other units a heuristic must be used to decide whether they must be changed.
Generally, the weights are determined by adjusting a set of initial weights by adapting them to a sequence of input_states and the associated target_states. The function learn is used for this purpose.

```
learn :: weights_net -> [input_state] -> [target_state] -> weights_net
learn weights (sample:samples) (target:targets)
    = learn (adapt_weights weights sample target) samples targets
learn weights [][]
= weights
```


## The generalized delta rule

The generalized delta rule is a supervised learning rule used, in one variant or another, in many networks. It will be proven that this rule 'works'; the weight changes prescribed by this rule reduce the error made by the network [Rumelhart 86]. The delta rule is used for networks containing McCulloch-Pitts neurons with a synapse that take a weighted sum. The proof is only valid for neurons with a nondecreasing differentiable generator function, but the delta rule is used with success for networks processing binary values. Unfortunately,
the only way to check whether the network has learned to execute its task is to test it with all possible input patterns. Another problem is that it is possible to get stuck in a sub-optimal solution.

In informal terms the ordinary delta rule formulated for a single layer of binary elements, devised in 1959, can be expressed as [Widrow-Hoff]:

- Present a sample to the network and compute the corresponding output;
- If the output is equal to the target do nothing;
- If the output value is high and the target value is low, the potential (the value of the inner product) is decreased by enlarging the weights corresponding to negative inputs and decreasing the weights associated with positive inputs;
- If the output is improperly low, the potential is increased by augmenting the contribution of positive inputs and lessening the influence of negative inputs.
The formula for changing the weights associated to input $i$ in element $j, W_{\mu}$, is

$$
\Delta W_{l l}=\alpha\left(\text { target }_{j}-\text { output }_{j}\right) \text { nput }_{l}
$$

The factor $\alpha$, in this formula, is a small positive value determining the learning rate. This simple formula is only applicable for single layered networks, since the target values for hidden elements are generally unknown. Moreover, it is not proven that this rule yields a correctly behaving network.

To derive the generalised delta rule, a measure for the (in)correct behaviour of the network is needed. The error for a specific pair of input and associated target state is half the square of the length of the difference vector of target and output.

$$
\begin{equation*}
E=\frac{1}{2} \sum_{t}\left(\text { target }_{t}-\text { output }_{t}\right)^{2} \tag{1}
\end{equation*}
$$

The overall measure of error is the summation of this error over all patterns; $E_{\text {total }}=\sum E$. The generalized delta rule implements a gradient descent in $E$. The weights are changed to reduce the error, $E$, by

$$
\begin{equation*}
\Delta W_{ر l}=-\alpha \frac{\partial E}{\partial W_{ر}} \tag{2}
\end{equation*}
$$

In the remainder of this derivation it is assumed that the synapses take a weighted sum (wsum), and that the generator is a nondecreasing differentiable function, $f$, of the current potential. This assumption makes this derivation inapplicable to all nets with binary values. The potential of element $j, v_{j}$, is given by

$$
\begin{equation*}
v_{J}=\sum_{t} W_{\mu} \text { input }_{l} \tag{3}
\end{equation*}
$$

The (output) value of this element is determined by

$$
\begin{equation*}
\text { output }_{j}=f\left(v_{j}\right) \tag{4}
\end{equation*}
$$

It is useful to write the derivative in 2 as the product of the change of the error as function of the potential and the change of potential as function of the weight.

$$
\begin{equation*}
\frac{\partial E}{\partial W_{\jmath \jmath}}=\frac{\partial E}{\partial v_{\jmath}} \frac{\partial v_{ر}}{\partial W_{ر}} \tag{5}
\end{equation*}
$$

From equation 3 it is obvious that the last term can be written as

$$
\begin{equation*}
\frac{\partial v_{1}}{\partial W_{\jmath t}}=\frac{\partial}{\partial W_{\jmath t}} \sum_{k} W_{j k} \text { input }_{k}=i n p u t_{t} \tag{6}
\end{equation*}
$$

Now the function $\delta$ is defined. This function gives the learning algorithm its name.

$$
\begin{equation*}
\delta_{j}=-\frac{\partial E}{\partial v_{j}} \tag{7}
\end{equation*}
$$

Using this definition, equation 5 can be rewritten as

$$
\begin{equation*}
\frac{\partial E}{\partial W_{\jmath 1}}=-\delta_{j} \text { input }_{l} \tag{8}
\end{equation*}
$$

An equivalent form for the weight changes according to equation 2 is

$$
\begin{equation*}
\Delta W_{\mu}=\alpha \delta_{j} \text { input }_{l} \tag{9}
\end{equation*}
$$

This leaves the task to determine $\delta$ for each node in the network. Using the rule of partial derivatives again, the definition is split into a product of error change resulting from variations in the output and output changes due to variations in the potential:

$$
\begin{equation*}
\delta_{j}=-\frac{\partial E}{\partial v_{j}}=-\frac{\partial E}{\partial o u t p u t_{j}} \frac{\partial \text { output }_{j}}{\partial v_{J}} \tag{1}
\end{equation*}
$$

From 4 it is clear that the last factor can be rewritten as

$$
\begin{equation*}
\frac{\partial \text { output }_{J}}{\partial v_{J}}=f^{\prime \prime}\left(v_{J}\right) \tag{11}
\end{equation*}
$$

For an output unit, the first term obtained in 10 is easily deduced from 1

$$
\begin{equation*}
\frac{\partial E}{\partial o u t p u t_{j}}=\frac{\partial}{\partial o u t p u t_{j}}\left(\frac{1}{2} \sum_{1}\left(\text { target }_{1}-\text { output }_{1}\right)^{2}\right)=-\left(\text { target }_{j}-\text { output }_{j}\right) \tag{12}
\end{equation*}
$$

By substituting these results, equation 10 becomes

$$
\begin{equation*}
\delta_{j}=\left(\text { target }_{j}-\text { output }\right) f^{\prime}\left(v_{j}\right) \tag{13}
\end{equation*}
$$

For a single layer of linear elements (the generator function is the identity function $f v=v$ ) the ordinary delta rule is obtained.

$$
\begin{equation*}
\Delta W_{\mu}=\alpha\left(\text { target }_{j}-\text { output }_{j}\right) \text { input }_{t} \tag{14}
\end{equation*}
$$

The chain rule is used to derive an expression for $\delta$ of hidden units. First, the leading term in equation 10 is rewritten

$$
\begin{align*}
\frac{\partial E}{\partial \text { output }_{j}} & =\sum_{k} \frac{\partial E}{\partial v_{k}} \frac{\partial v_{k}}{\partial \text { output }_{j}}=\sum_{k} \frac{\partial E}{\partial v_{k}} \frac{\partial}{\partial \text { output }_{j}} \sum_{l} W_{k t} \text { input }_{l} \\
& =\sum_{k} \frac{\partial E}{\partial v_{k}} W_{k j}=-\sum_{k} \delta_{k} W_{k J} \tag{15}
\end{align*}
$$

The summation is over all units receiving output from element $j$. Using this result and equation $I I$ to rewrite $10, \delta$ becomes

$$
\begin{equation*}
\delta_{j}=f^{\prime}\left(v_{j}\right) \sum_{k} \delta_{k} W_{k j} \tag{16}
\end{equation*}
$$

Using equations 13 and $16, \delta$ can be computed for every unit in the network. The error measure, $\delta$, is propagated from the output units backwards. This gives rise to the name backpropagation for this algorithm. With these values, the weights can be adapted, according to 9 , to reduce the error.

When the network behaves well for all possible inputs, it will do forever; the weights are only changed when an erroneous output is generated. From 9 it can be deduced that $\delta$ must be non zero to change the weights, but from 13 and 15 it is clear that all $\delta$ 's are zero when the network produces the correct output.

It is easy to show how a threshold value involved in the generation of the output can be adapted to reduce the error. Assume that the output is determined by

$$
\begin{equation*}
\text { output }_{j}=f\left(v_{j}+t_{j}\right) \tag{17}
\end{equation*}
$$

The error is reduced, in a gradient descent manner, by changing the threshold, $t$, by an amount

$$
\begin{equation*}
\Delta t_{j}=-\beta \frac{\partial E}{\partial t_{j}} \tag{18}
\end{equation*}
$$

Analogue to the derivation above, it can be shown that

$$
\begin{equation*}
\Delta t_{j}=\beta \delta_{j} \tag{19}
\end{equation*}
$$

Obviously, the threshold can be replaced by an additional input connected to a constant element with value $\beta / \alpha$. Usually the value 1 is used for $\beta / \alpha$.

Changing the weights after each presentation of an input pattern and associated target pattern departs slightly from the true gradient descent for $E_{\text {total }}$.

Provided that the learning rate, $\alpha$, is sufficiently small, the delta rule will implement a very close approximation to the gradient descent in sum-squared error. The rule will search for the least mean square error, hence this algorithm is sometimes named the LMS-rule.

The algorithm shown here is a gradient descent method and cannot distinguish between the global minimum of the error function in the weights space and local minima. This problem resembles the problem of avoiding local maxima in any hill climbing algorithm. Fortunately, simulation results show that local minima are usually avoided.

Even when local minima are avoided, there is a problem of convergence speed. Clearly the factor $\alpha$ in equation 9 above, cannot be made very large to maintain the gradient descent character. A true gradient descent procedure requires even that infinitesimal steps are taken. In the discussion of networks below we will mention the learning rate and number of needed samples to teach the network some simple functions. There are also extensions proposed, to the learning algorithm as presented above, which should increase the learning speed; e.g., Rumelhart [Rumelhart 86] proposes the introduction of a momentum term in the change of weights. Recent theoretical work suggests that better results will be obtained by adapting the weights after each presentation of a pattern by a relatively large amount [Ellacott 90].

The learning algorithm presented here has another problem. When all initial weights are chosen identical in a symmetrical network and the problem desires unequal weights, the procedure will never find these weights. The error propagated back through the network is proportional to the weights: if all weights are equal all $\delta$ s will be equal. Due to the symmetry, also the outputs of the elements are equal, these outputs are used as inputs by the next layer. Equation 9 yields identical changes of the weights for all units and the symmetry is preserved forever. This problem can easily be avoided by choosing some small pseudo random values for the initial weights, or adding a little noise somewhere in the system or the learning algorithm.

The final problem is that the derivation is only valid for differentiable and nondecreasing generator functions; it not applicable in binary networks. This learning algorithm is specified in Miranda when it is used in the multi-layer perceptron described below.

In binary networks a similar learning algorithm is used. The derivation terms are simply dropped. It cannot be proven to work, but it is used with good results in many applications. This algorithm is described in detail in the discussion of multi-layer perceptron for binary values below.

## Description of specific artificial neural networks

With the description tools defined above and some additional, model specific, functions it is possible to give an elegant, complete and executable description of artificial neural networks. For each network at least the following functions must be defined:

| high | $::$ value | II The upper limit of the values |
| :--- | :--- | :--- |
| low | $::$ value | II The lower limit of the values |

### 6.2 The Adaptive Linear Combiner

The ALC (Adaptive Linear Combiner) is the basic building stone for many adaptive systems. In terms of neural networks, the ALC is only a synapse that takes a weighted sum. It is developed independently in adaptive signal processing [Widrow 85].

$$
\begin{aligned}
& \text { alc :: element } \\
& \text { alc }=\text { wsum }
\end{aligned}
$$

With its input connected to a tapped delay line, the ALC is the most popular key component of an adaptive filter. When the ALC is used as an adaptive filter, an estimation of the wanted output must be available to compute an error signal. This error is the difference between the system output and the desired output. The adaptive filter is tuned to reduce the error signal. There are several algorithms to decrease the error signal, the most popular method is the delta rule.

The use of these adaptive filters is widespread nowadays; all high speed modems use adaptive equalization filters, many long distance communication links are equipped with adaptive echo cancelers, furthermore these filters are used in noise cancelers and system modeling.

Since our prime interest is the description of neural networks, the ALC is not discussed in detail. It is mentioned to show that even very simple elements can perform useful tasks.

### 6.3 The Adaline

The name adaline is derived from adaptive learning element [Widrow 62]. Adalines are adaptive linear combiners cascaded with the sgn function to get a binary output [Widrow 88].

```
adaline :: eiement
adaline weights = sgn.wsum weights
```

An adaline computes the inner product between its weight vector and the input vector and passes this value, the potential, to the sgn function to obtain a binary output. Usually, a constant value is added to the input state to obtain an adjustable threshold.

The output of an adaline changes when the inner product crosses zero. The inner product of the input and the weight vector is zero in a plane through the origin perpendicular to the weight vector. So, the adaline checks whether the input is above or below a hyperplane in the input space. The regions in the input domain correspond to output values. Binary functions with an input domain that can be separated in two regions, separated by a hyperplane, can be performed by the adaline. This implies that an adaline can learn to execute on logical AND or OR-function, but the XOR-function cannot be performed. The typical adaline 'network' contains exactly one element.

```
compute_state :: weights_net -> input_state -> net_state
compute_state = feed_forward adaline
compute_output }\because\mathrm{ : weights_net -> input_state -> output_state
compute_outpul weights inputs = select_output (compute_state weights inputs)
initial_net :: [num] -> weights_net
initial_net [dımension,elements]
= [[[a*(1-((2*1+\mp@subsup{3}{}{*}])\operatorname{mod}7)/3) || Pseudo random we1ghts
    | | <- [1..dimension+nr_constant_neurons]] || Length of input vector
    | J<-[1 elements]]] | The number of elements
high = 1 II The standard convention Moreover,
low = -1 |l it is assumed that the mean input is 0.
nr_constant_neurons = 1
```

The weight vectors for some familiar boolean functions are (the last element represents the threshold value)

```
and_vector = [1,1,-15] |l 2-mput AND
nand_vector = [-1,-1, 1 5] || 2-mput NAND
or_vector = [1,1, 1.5] |l 2-mput OR
maj_vector = [1,1,1,0] |l 3-mput majonty
```

Our specification shows that an adaline is just a special case of the perceptron. So, learning in adalines is equivalent to learning in perceptrons and it is described there.

In the early 1960s an extension to implement non linear separable functions is proposed [Ridgway 62], called madaline (many adalines). These networks consist of a layer of adalines connected to a single, fixed, logic device in the second layer. It can be modelled by feeding the output from the network described above to the proper function. Some useful functions are

```
and_fun :: input_state -> output_state
and_fun inputs
    = [high], if # (filter (=low) inputs) =0
    = [low], otherwise
    majority_fun :: input_state -> output_state
majority_fun inputs
    = [high], if # (filter (=high) inputs) > # inputs/2
    = [low] . Otherwise
```

A possible madaline performing the 2 -input XOR-function is

```
xor_madaline = and_fun.compute_output [[or_vector, nand_vector]]
```


### 6.4 The Perceptron

Perceptrons [Rosenblatt 62, Minsky 69] are feed forward networks composed of simple processing elements. These processing elements are similar to the adaline described above. In the original setting, perceptrons are processing binary values. Later on a variant to process continuous variables is defined in order to apply the generalized delta rule as learning algorithm. First, the single layer binary perceptron will be described, afterwards multi-layer perceptrons for continuous and binary values will be discussed. It makes sense to distinguish single-layer and multi-layer perceptrons; the type of mappings which can be executed by these networks and the applicable learning rules are different.

The internal potential in each element of a binary perceptron is determined by the function wsum; the inner product between the input and weight vector. The generator is the hard_limit function. Usually, the threshold in this function is taken to be zero. An adaptable threshold is modelled by an additional constant input. When the bounds high and low are set to 1 and -1 , which is the common choice, the single layer perceptron consists of adalines.

## The single layer perceptron

In the literature, the single layer perceptron is sometimes drawn as a three layer network; a layer of input neurons yielding the input state, a layer of unalterable elements performing the fixed mapping from the input to the adaptable elements and the layer of adaptable elements. The description given here covers only the last layer, the only layer that is adapted.

The single layer perceptron contains one neuron for each element in the output state. The functional description of the single layer perceptron is identical to the adaline as described above, only the element used in compute_state is changed from adaline to perceptron_b. Binary functions with an input domain that
can be separated in two regions corresponding to the output values (like the AND and OR-function) can be implemented by a single-layer perceptron.

There are effective ways to teach a perceptron to execute a function based on a set of sample inputs and the desired response. Since the generator function is not differentiable the derivation of the learning rule found as generalised delta rule is not valid. However, the normal delta rule is used to adapt the weights in this network after every presentation of a sample. This makes sense, since the generator function can be thought to magnify the error, the delta rule will improve the potentials in the network and in the long run also the output values. The correctness of this leaming rule is not proven, but it appears to behave well in simulations. A set of small randomly chosen weights is chosen for the initial network. After each presentation of a sample the weights are adapted such that the error in the response of the network is decreased. The Miranda specification of the learning algorithm for the single layer binary perceptron reads:

```
adapt_weights:: weights_net -> input_state -> target_state -> weights_net
adapt_weights weights input target
= [[[new_weight ji
    | i<- [0..# weightsII!j -1]] |l inputs for element i in the layer
    | << [0.# weights!! -1]]] |l elements in the layer
    where new_weight ji = weights!0!!!i + gain * state!0!i * error!j
        state = compute_state weights input
        output = select_ouput state
        error = [ target!i - output!i|i<-[0..# target -1]]
```

Functions with an input domain that cannot be separated by a single hyperplane (like the XOR-function), cannot be implemented by a single layer perceptron. More complicated functions must be implemented by identifying smaller regions in input space. There are at least two possibilities to implement these functions with neural networks. First, the domains in the input space corresponding to each output value can be limited by several hyperplanes, each plane can be implemented by an adaline. The information yielded by these elements can be further processed in other layers of these elements. It can be proven that these multi-layer perceptrons can learn every projection function. The second way to implement these functions is to use neurons sensitive to smaller regions in the input space. The Learning Vector Quantization network uses a set of vectors, every input vector is associated with the vector which resembles this input most.

## Applications of the single layer perceptron

This single layer perceptron network learns simple two input functions like AND and OR within some tens of samples, the exact number of samples needed depends on the initial weights, the gain and the sequence of examples.

These functions are also learned when some don't-care elements are added to the input state. The weights associated with these don't-care elements become
zero (or very small) as expected. Using gain $=0.1$ and a scattering of initial weights $a=0.1$, this net learns the two input AND-function from 20 randomly chosen input patterns. The resulting weight vector of a sample run is [027, 0.10,0 23]. The last element in this vector is the connection to the constant input, representing the threshold. A three input AND-function takes 80 samples, the resulting vector is [027, 0 30, 0 37,-0.77]. When a don't-care term is added (the target for the function trace in appandix C is given by the function and_fun take 2) the resulting vector is $[0.47,030,-0.03,-0.37]$. This requires 30 leaming samples. The weight vectors obtained are dependent of the initial weights and the sequence of samples.

## The multi-layer perceptron for continuous values

For the multi-layer perceptron, with a differentiable generator function, the generalized delta rule is used as learning algorithm. The generator function is the sigmoid defined above. This network can be specified in Miranda as:

```
compute_output :: weights_net -> input_state >> output_state
compute_output weights input = select_output (compute_state weights input)
compute_state :: weights_net -> input_state -> net_state
compute_state = feed_forward perceptron_c
adapt_weights:: weights_net -> input_state -> target_state ->weights_net
adapt_weights weights input target
= [[[new_weight l ] I
    | i <- [0..# weights'I'j -1]] II inputs for neuron I in layer I
    ||<-[0..# weights!l -1]] Il neurons in layer I
    | <- [0 # weights -1]] Il layers
    where new_weight l | | = weights||| + gain * statelll| deltall
    state = compute_state weights input
    state'
        = [[ sıgmord' (wsum (weights'll|) (statell)) ||<< [0. #weights|| -1]]
                            | l<-[0 #weights -1]]
    output = select_output state
    output' = ast state'
    error = [output'J*(target'J - output'J)| | <- [0 # target -1]]
    delta
            = [[ state'|l] * sum [werghts'(l+1)'k}\mp@subsup{|}{}{\prime}\mp@subsup{}{}{*}\mp@subsup{\operatorname{delta}}{}{\prime}(l+1)|
            |k<-[0..# werghts'(l+1)-1]| || outputs of neuron
            | |<-[0 # weights| -1]] |l neurons in layer I
            | <- [0. # weights -2] ] ++ |l hidden layers
            [error] Il output layer
```

```
initial_net :: [num] -> weights_net
initial_net (n:m:sizes)
= [layer] \(\quad\), if sizes \(=[\) ]
= layer : initial_net (m:sizes) , otherwise
    where layer
\(\begin{aligned}= & {\left[\left[a^{*}\left(1-\left(\left(613^{*} \# \text { sizes }+337^{*} i+733^{*} j\right) \bmod 101\right) / 50\right) \quad \text { I| pseudo random }\right.\right.} \\ & \left.\left.\mid j<-\left[1 . . n+n r^{\prime} \text { constant_neurons }\right]\right] \mid i<-[1 . . m]\right]\end{aligned}\)
```

high $=1$
low $=-1$
nr_constant_neurons = 1

## Application of the multi-layer perceptron for continuous values

The behaviour of this network is studied, it was trained to execute some boolean functions. Using boolean functions there is a problem in deciding when the network has learned the function. Due to the used generator function the extremes high and low are never reached as output values. A large sharpness improves the separation between the output values, but tends to slow down the recovery from wrong situations since the derivative tends to zero very soon. We have used as parameters $a=0.2$, gain $=0.5$ and sharpness $=2$. This results in a quick convergence for the three input majority function, within 100 samples the output values are within $10 \%$ of the difference between high and low of the correct values. The initial vector with value $[-0.04,-0.14,0.16,0.05]$ is changed to $[1.47,1.46,1,45,0.14]$. A two layer network with two hidden units and one output unit learns the XORfunction within 190 samples. The initial weights generated by initial_net $[2,2,1]$ are [[[ $[0.08,-0.02,-0.02],[-0.05,-0.16,0.14]],[[-0.04,-0.14,0.16]] 1$. The obtained weights are [[[1.21,-1.25,-0.90], [1.85,-1.56,1.85]], [[1.71,-1,67,1.47]]]. To show how these vectors implement the XOR-function the truth tables for the processing elements are shown; the output values are rounded to boolean values. Elements 1 and 2 are hidden, element 3 is the output neuron.

| inputs |  | element 1 | element 2 | element 3 | system |
| :--- | :--- | :--- | :--- | :--- | :--- |
| low | target |  |  |  |  |
| low | low | low | high | high | low |
| high | low | low | low |  |  |
| high | low | high | high | high | high |
| high | high | high |  |  |  |
| high | low | high | high | low | low |

Multi-layer perceptrons with the backpropagation learning rule are used with success in many practical applications [see for instance INNC 90]. The main problem is to develop a suitable mapping from the problem variables to the neural network inputs and to use the network output in the appropriate way. The number of hidden units needed, the gain and the initial network are determined empirically.

## The multi-layer perceptron for binary values

One way to ensure that the output reaches the bounds high and low is to use a binary generator function. Again, there is no learning rule which can be proven to reduce the error, but a generalisation of the single layer learning rule for binary values, the delta rule, can be used here. Compared with the generalized delta rule, this comes down to dropping the derivation terms. The single layer perceptron is just a special case of this network. This network differs from the network for continuous values in the element used to compute the state and the learning rule:

```
compute_state :: weights_net -> input_state -> net_state
compute_state \(=\) feed_forward perceptron_b
adapt_weights \(::\) weights_net -> input_state -> target_state ->weights_net
adapt_weights weights input target
= [[[new_weight I j i
    | \(\mathrm{i}<-[0 . . \#\) weights!!! \(\mathrm{j}-1]] \quad\) II inputs for neuron j in layer l
    \(\mid \mathrm{j}<-[0 . . \#\) weights \(\mid 1-1]] \quad\) I| neurons in layer 1
    \(\mid 1<-[0 . . \#\) weights -1\(]] \quad\) || layers
    where new_weight \(1 \mathrm{j} i=\) weightsIlljli + gain * state!lli * deltalllj
        state \(=\) compute_state weights input
        output = select_output state
        error \(=\) [target! j - outputlj \(\mid j<-[0 . . \#\) target -1\(]\) ]
        delta
        \(=\left[\mid\right.\) sum [weights! \((1+1) \mid k!j{ }^{*}\) delta! (l+1)]k
            \(\mid k<-[0 .\). \# weights!( \(1+1\) ) -1]] II outputs of neuron
            | \(\mathrm{j}<-\) [0..\# weights! 1 -1] ] || neurons in layer l
            | \(<\) <- [0..\# weights -2] ] ++ || hidden layers
            [error] I| output layer: error
```

A three layer perceptron can perform every projection of the input state to the output state provided there are enough elements in the layers [Minsky \& Papert 69]. The first two layers select the proper regions in the input space while the third layer combines the output of these combinations. Unfortunately, the amount of processing elements required cannot be determined.

There is a very crude implementation of every boolean function on a twolayer perceptron [Coolen 89]. Its use is not recommended, but it shows the power of multi-layer perceptrons. The implementation of the function for each output bit is done with brute force:

- For every possible input pattern that corresponds to a high output value there is a hidden neuron that selectively responds to that pattern. This is done by making the weight vector identical to the input pattern and choosing the proper threshold (square of the length).
- The output bit can now be determined by a perceptron taking the logical OR of these hidden units.


## Application of the multi-layer perceptron for binary values

Using this network description, the number of training samples needed to learn the XOR-function as function of the amplitude of initial weights, a, the learning rate, gain, and the number of neurons in the hidden layer is explored. The XORfunction is used because it is a very simple and small function, but still interesting through a non linear separable input-output releation. Experiments were done with the same pseudo random sequence of samples and the same set of pseudo random initial weights (except for the amplitude). To study the influence of the particular choice of initial weights some additional experiments were done; a negative amplitude inverts all initial weights and a different set of initial weights is used as additional measurement, these results are printed in italics. After several learning samples the performance of the system was evaluated until the XOR-function was learned. The number of samples between the subsequent evaluations of the network performance is always less then $10 \%$ of the total number of samples needed. The Miranda system on a SUN-3 performs around 200 updates of individual connection strength per second, including evaluation of the performance and printing statistics. Our patience and the amount of heap space used limit the number of learning cycles somewhere around 10000 . The number of samples corresponding to the optimal gain for a given amplitude is printed bold.

| $\downarrow$ a $\mid$ gain $\rightarrow$ | 0.0005 | 0.001 | 0.002 | 0.005 | 0.01 | 0.02 | 0.05 | 0.1 | 0.2 | 0.5 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.0001 | 150 | 40 | 350 | 3100 | 2050 | 1350 | $>10000$ | $>10000>10000$ |  |  |
| 0.001 | 1450 | 480 | 100 | 270 | 5270 | 4575 | $>10000$ | $>10000$ | $>10000$ |  |
| -0.001 | 1400 | 525 | 125 | 325 | $>10000$ | 675 | 750 | 100 | $>10000$ |  |
| 0.01 | 9850 | 3200 | 850 | 100 | 50 | 350 | 3175 | 950 | 3400 | $>10000$ |
| 0.01 | $>10000$ | 3175 | 800 | 225 | 40 | 2875 | 7875 | 8100 | $>10000$ | $>10000$ |
| -0.01 | $>10000$ | 3300 | 875 | 120 | 60 | 170 | 5750 | 7350 | 1375 | $>10000$ |
| 0.1 | $>10000$ | $>10000$ | 9525 | 975 | 300 | 60 | 6000 | 800 | 425 | 3300 |
| 1 | 2850 | 1550 | 790 | 280 | 150 | 80 | 55 | 30 | $>10000$ | $>10000$ |

Number of teaching samples needed by a 2 -layer binary perceptron to leam the 2 -input XOR-function with 2 elements in the hidden layer.


The data from the table above shown graphically. Vertical the number of samples needed, horizontal the gain.

| $\downarrow$ algain $\rightarrow$ | 0.0005 | 0.001 | 0.002 | 0.005 | 0.01 | 0.02 | 0.05 | 0.1 | 0.2 | 0.5 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.0001 | 95 | 925 | 800 | 5625 | 175 | 3800 | 2500 | $>10000$ |  |  |
| 0.001 | 1475 | 310 | 70 | 150 | 550 | 350 | 6650 | $>10000$ |  |  |
| 0.01 | $>10000$ | 3675 | 1900 | 80 | 45 | 1800 | 6625 | 650 | 6950 | $>10000$ |
| 0.1 | $>10000$ | $>10000$ | $>10000$ | 1825 | 450 | 210 | 40 | 3200 | 1900 | 1500 |

Number of teaching samples needed by a 2 -layer binary perceptron to learn the 2 -mput XOR-function with 3 elements in the hidden layer.

| Lalgain $\rightarrow$ | 0.0005 | 0.001 | 0.002 | 0.005 | 0.01 | 0.02 | 0.05 | 0.1 | 0.2 | 0.5 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.0001 | 150 | 900 | $\mathbf{8 0}$ | 2675 | 1250 | 400 | 850 | 750 | $>10000$ |  |
| 0.001 | 875 | 510 | $\mathbf{1 2 5}$ | 1240 | 5600 | 1250 | 850 | 750 | $>10000$ |  |
| 0.01 | $\mathbf{9 8 0 0}$ | $\mathbf{2 8 7}$ | $\mathbf{5 2 5}$ | 100 | $\mathbf{4 5}$ | $\mathbf{3 7 7 5}$ | 5350 | 1725 | 4850 | $\mathbf{8 7 2 5}$ |
| 0.1 | 10000 | $>10000$ | 5025 | 1125 | $\mathbf{3 7 5}$ | $\mathbf{1 6 0}$ | 225 | $\mathbf{3 0 0}$ | 3400 | 3800 |

Number of teaching samples needed by a 2 -layer binary perceptron to leam the 2 -input XOR-function with 4 elements in the hidden layer.

| $\downarrow_{\text {a }}$ gain $\rightarrow$ | 0.0005 | 0.001 | 0.002 | 0.005 | 0.01 | 0.02 | 0.05 | 0.1 | 0.2 | 0.5 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.001 | $>9400$ | $>9400$ | $>9400$ | $>9400$ | $>9400$ | 4150 | 3700 | 4750 | 1700 | $>10000$ |
| 0.01 | $>9400$ | $>9400$ | $>9400$ | $>9400$ | 6375 | $>9425$ | 2100 | 1800 | 3100 | 1700 |

Number of teaching samples needed by a 2 -layer binary perceptron to leam the 3 unput XOR-function with 2 elements in the hidden layer.

| 】a\|gain $\rightarrow$ | 0.0005 | 0.001 | 0.002 | 0.005 | 0.01 | 0.02 | 0.05 | 0.1 | 0.2 | 0.5 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.001 | 1225 | 300 | $\mathbf{1 5 0}$ | 800 | 3250 | 6100 | 7450 | $>10000$ | 700 | $>10000$ |
| 0.01 | $>9175$ | $\mathbf{3 1 0 0}$ | $\mathbf{8 0 0}$ | $>9125$ | $\mathbf{8 5}$ | $>9050$ | $\mathbf{3 6 0 0}$ | 700 | 600 | 1600 |

Number of teaching samples needed by a 2 -layer binary perceptron to leam the 3 input XOR-function with 3 elements in the hidden layer.

From these figures some tendencies can be observed:

- The 2-input XOR-function is quickly learned for each number of hidden units tried, provided that the initial parameters are well chosen.
- Amplitude and gain can be chosen in a wide range of values to achieve a correctly behaving network. There is however, a clear optimum gain for each amplitude. This optimum is largely independent from the actual initial weights; the amplitude of these weights is much more important. The optimum gain increases when the amplitude of the initial weights increases.
- The 3 -input XOR-function appears to be much more difficult than the 2 input XOR-function. The 3 -input network behaves nearly correct (only one pattern gives the wrong answer) within some hundreds of samples, but it requires much time to become totally correct. The behaviour of the network keeps changing in time, so it is not stuck in a local minimum of the error function.
The internal representations of the XOR-function can be deduced from the obtained weights. There is not a unique internal representation for a given network, but the representation changes among the learn runs.


### 6.5 Learning Vector Quantization

The Learning Vector Quantization (LVQ) network [Kohonen 88] consists of a single layer of elements. All elements determine the correspondence between the input vector and their weight vector by computing the inner product. The output of the network is the class associated with the element with the largest potential. When several elements produce an equal potential a choice is made; here the element with the lowest index number is used. The element with the largest potential can be determined in the neural network manner by a competition of elements in an N -flop feedback network (see below), or in a more direct way. Here the direct way is used; the maximum potential is determined and the winner is the first element with that potential.

During the learning phase the winning weight vector is adapted; when the output class is equal to the target the weight vector is changed in the direction of the input to increase the inner product further. In case of an erroneous output, the winning weight vector is changed in the opposite direction to decrease the potential of the winning element.

There is no need to use random vectors as initial values for the weights in this network. It appears that the best results are obtained when all initial vectors are chosen in the middle of the input space.

Unlike learning according to the delta rule, this network is adapted when it behaves correctly for the current input-target pair. As a consequence, it is possible that the network performance decreases even if it had learned to execute its task perfectly. This can be understood by considering a vector responsible for the correct mapping of inputs in some region in the input space, a sequence of inputs on one side of this region drifts the weight vector to this
side. The other side of the region becomes closer to another weight vector and hence, possible to another output class. This happens only when a single weight vector maps several different inputs to their targets. This effect can be reduced by slowly decreasing the learning rate.

| vector | $==$ | [num] |
| :--- | :--- | :--- |
| vector_id | $==$ | num |
| class | $==$ | state |
| target_class | $==$ | class |
| nr_off_classes | $==$ | num |

associate_class :: vector_id -> class
associate_class id = [id mod nr_off _classes] || equal distribution over classes
compute_state $::$ [vectors] $>$ input_state $->$ net_state
compute_state $=$ feed_forward inner_product
compute_output :: [vectors] -> input_state -> output_state
compute_output vectors input = associate_class (find_winner vectors input)
find_winner $::$ [vectors] -> input_state -> vector_id
find_winner vectors input = id_max (select_output (compute_state vectors input))
|| more direct: id_max [inner_product vector input | vector <- vectors]
id_max :: output_state $->$ vector_id
id_max state = scan_list 0 (max state) state
adapt_weights :: [vectors] -> input_state $->$ target_class $->$ [vectors]
adapt_weights [vectors] input target
= [[new_vector $\mathrm{i} \mid \mathrm{i}<-[0 . . \#$ vectors -1$]]]$
where
new_vector i
$=$ vectorsli , if i = winner
$=$ [vectorslilj + gain * input! $\mid \mathrm{j}<-$ elements] , if associate_class winner = target
$=$ [vectorslilj - gain *inputlj $\mid j<-$ elements ], otherwise
elements $=$ [0..\# vectorsli -1]
winner = find_winner [vectors] input
initial_net :: [num] -> [vectors]
initial_net [input_dimension,size]
$=$ [[vector $i \mid i<-[1 . . s i z e]]]$
where vectori
$=\left[(\right.$ high+low $) / 2+a^{*}\left(\left(2^{*} j+3^{* i}\right) \bmod 7\right) / 6$
| $\mathrm{j}<-$ [1..input_dimension + nr_constant_neurons] ]

```
high \(=1 \quad\) Il the usual assumption
low \(=-1\)
nr_off_classes = \(2 \quad\) |l binary output
```

```
nr_constant_neurons = 0
```

For vectors not distributed at a fixed distance around the ongin, the inner product is not applicable as a measure of correspondence. The length of the difference vector can be used as a measure of correspondence in this situation [Kohonen 88]. The element with the least potential must be the winner in this situation. In the learning algorithm described above the weight vectors can grow unbounded to resolve some severe errors of the network. The learning algorithm has been adapted to decrease the length of the difference vector when the network produces the correct answer.

```
compute_state :: [vectors] -> input_state -> net_state
compute_state = feed_forward difference
difference :: vector -> vector -> num
difference v1 v2 = length [v1!ı - v2! |i<< [0 #v1-1]]
compute_output :: [vectors] -> input_state -> output_state
compute_output veclors input = associate_class (find_winner vectors input)
find_winner :. [vectors] -> input_state >> vector_id
find_winner vectors input = id_min (select_oulput (compute_state vectors input))
id_min ` output_state -> vector_Id
id_min state = scan_list 0 (min state) state
adapt_weights :: [vectors] -> input_state -> target_class -> [vectors]
adapt_weights [vectors] input target
= [[new_vector ו | | <-[0. # vectors -1]]]
    where
    new_vector I
        lol
```

There are many similar networks proposed in the literature. The Hamming network [Lippmann 87] uses the Hamming distance between the input bit vector and stored vector to select the best matching vector. The Hamming distance between two bit vectors is the number of corresponding elements that are unequal. The counter propagation network [Hecht-Nielsen 87a] does not use a fixed association of classes to vectors, but this relation is learned in a perceptron like manner from the samples.

## Applications of Learning Vector Quantization

The LVQ network learns very quickly. Using the network description given above the number of teaching samples needed for various values of the initial values and gain is determined. The 2-input XOR-function was taught to a network consisting of 4 vectors. Several variants of the learning algorithm are compared. In the first table the original learning algorithm is used. The second table contains the results obtained using the second algorithm. In the last learning algorithm the vectors corresponding to the right classes are normalised. The performance of the last algorithm decreases very significantly if the wrong vectors are normalised after adaptation. The network becomes extremely sensitive for the choice of the initial values and the sequence of training inputs. The Miranda system on a SUN-3 performs around 30 vector updates per second.

| $\downarrow$ algain $\rightarrow$ | 0.0001 | 0.001 | 0.01 | 0.1 | 1 | 0.0001 | 0.001 | 0.01 | 0.1 | 1 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0 | 21 | 21 | 21 | 21 | 21 | 19 | 19 | 19 | 19 | 136 |
| 0.0001 | 18 | 18 | 18 | 11 | 11 | 11 | 11 | 11 | 11 | 3 |
| -0.0001 | 6 | 19 | 19 | 18 | 5 | 4 | 5 | 5 | 4 | 5 |
| 0.001 | 18 | 11 | 18 | 11 | 18 | 11 | 11 | 12 | 11 | 12 |
| -0.001 | 4 | 4 | 4 | 18 | 4 | 4 | 4 | 4 | 2 | 2 |
| 0.01 | 18 | 3 | 11 | 11 | 18 | 12 | 3 | 11 | 11 | 12 |
| -0.01 | 6 | 4 | 4 | 4 | 4 | 5 | 4 | 4 | 2 | 2 |
| 0.1 | 68 | 18 | 3 | 11 | 18 | 68 | 12 | 3 | 11 | 12 |
| -0.1 | 86 | 6 | 4 | 4 | 19 | 81 | 5 | 4 | 4 | 5 |
| 1 | 800 | 68 | 18 | 7 | 18 | 550 | 58 | 11 | 4 | 11 |
| -1 | 830 | 82 | 6 | 6 | 6 | 570 | 51 | 5 | 4 | 4 |

Left: Number of teaching inputs required by a 4 -vector leaming vector quantization network to leam the 2 -input XOR-function. The original learning rule is used.
Right: Initial values and trainings set as above. The second learning rule is used for this LVQ network, this rule yields vectors of bounded size.

| $\downarrow$ initial vectors g gain $\rightarrow$ | 0.0001 | 0.001 | 0.01 | 0.1 | 0.5 | 0.9999 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| all vectors [1,1] |  |  | $>8575$ | $>10000$ | 22 | 31 |
| all vectors [0,1] | 5 | 5 | 5 | 5 | 5 | 5 |
| pseudo random vectors | $>6500$ | 2 | 2 | 2 | 11 | 11 |
| inverse of vectors above |  | $>10000$ | $>10000$ | 6 | 4 | 11 |

The same samples are used again to teach this LVQ network. Now the vectors are normalised after a successful recognition in order to keep their length finite.

From these figures it is clear that the learning vector network learns the XORfunction much faster than a multi-layer perceptron using the backpropagation learning rule. Only when the scattering of initial vectors is much larger than the gain a large number of learning cycles is needed. The second learning rule keeps
the length of the vectors finite and performs slightly better than the others so it is recommended to use that one. It behaves much better when the distribution of samples over the vectors is not uniform. The normalising variant appears to be sensitive for the initial values and is not applicable when one of the vectors becomes the 0 -vector. This variant does behave well, when the bounds low and high are set to 0.1 and 1.0.

### 6.6 Feedback Networks

Networks with an internal loop structure, a connection of (part of) the outputs of some elements back to their inputs, are called feedback networks. All networks can be applied in a feedback configuration by making an external connection between the output and the input.

Feedback networks became famous from the work of Hopfield [Hopfield 82], although they were first used by Kohonen and Little [Kohonen 77, Little 74]. These networks consist of a feedback configuration of McCulloch-Pitts neurons. These elements take a weighted sum of their input values and yield usually a binary output value. When the weight vectors of all neurons are viewed as a matrix, the potentials are computed by multiplying this matrix with the input vector. The output vector is obtained by mapping the generator function over the potential vector. The state is cycled through the network until a stable state is reached.

The Hebb [Hebb 49] learning rule is used to tune this network. The strength of a connection is increased when the connected input and target value correspond, otherwise the weight is decreased. The connections from a neuron to its own input are kept zero in the historical setting, although the network behaves well with a non zero diagonal term in the weight matrix. Using this learning rule the network becomes an associative memory. The network can be used for pattern recognition. The proper weights are usually obtained by using the Hebb rule with each pattern used simultaneously as input and target. A pattern is also learned from multiple presentations, each one polluted with random noise. The patterns become the eigen vectors of the matrix multiplication cascaded with the generator function. The network seeks the pattern corresponding to an input state by re-applying the function until a fixed point is reached. The maximum number of patterns that can be stored in a large weight matrix is $14 \%$ of the number of neurons, provided that the patterns are not too similar [Hopfield 82, Amit 85].

Simulations show that the network finds the patterns faster when the neurons are updated one by one. Each neuron uses the most recent state to determine its output and unstable equilibrium points are avoided by the pseudo random evaluation order.

The bounds low and high are chosen nowadays usually -1 and 1, but Hopfield used 0 as lower bound. This symmetry has as result that the inverse of a learned pattern is also a fixed point of the system.

```
compute_state :: weights_net -> input_state -> net_state
compute_state = feed_back_asyn hopfield_b
compute_oulput \(::\) weights_net \(->\) input_state \(\rightarrow\) output_state
compute_output weights input = select_output (compute_state weights input)
adapt_weights:: weights_net -> input_state -> target_state ->weights_net
adapt_weights weights_net input_state target_state
\(=[[[\) new_weight \(i j \mid j<-[0 . . n-1]] \mid i<-[0 . . n-1]]]\)
    where new_weightij
\(\begin{aligned} & =0 \\ & =\text { weights_netlOlilj + gain * target_stateli * input_statelj, otherwise } \\ n & =\# \text { (weights_nello) }\end{aligned}\)
initial_net :: [num] -> weights_net
initial_net [ n ]
\(=\left[\left[\left[a^{*}\left(1-\left(\left(2^{*} i+3^{*} j\right) \bmod 7\right) / 3\right) \mid i<-[1 . . m]\right] \mid j<-[1 . . m]\right]\right] \quad\) II pseudo random
    where \(\mathbf{m}=\mathrm{n}+\mathrm{nr}\) _constant_neurons \(\quad\) II square matrix
```

```
high \(=1\)
```

high $=1$
low $=-1$
low $=-1$
nr_constant_neurons $=0$

```
nr_constant_neurons \(=0\)
```

The system appears to be sensitive for transformations of the input state. A translated pattern is not recognized. This can be partly solved by learning the network the inverse of the transformation expected. A large number of random patterns is used as sample while the target is the transformed pattern. It is of course, possible and more efficient to start with a matrix that performs the desired transformation. For transformations it is clearly better to use a synchronous update of the neurons, here an ordinary matrix multiplication is desired. When transformation and pattern recognition are taught in the right ratio, the network transforms an input pattern until it is recognized.

## Convergence

In order to show the convergence of the state in asynchronously updated binary feedback networks, an energy is associated with each state $s$ of the system. It will be shown that stable states are energy minima and that the system will evolve to a minimum. The energy $E$ is defined by

$$
\begin{equation*}
E=-\frac{l}{2} \sum_{i j} s_{i} w_{i j} s_{j} \tag{I}
\end{equation*}
$$

The output for element $i$ is given by

$$
\begin{equation*}
s_{i}^{\prime}=\operatorname{sgn}\left(v_{i}\right) \tag{2}
\end{equation*}
$$

$$
\begin{equation*}
v_{t}=\sum_{J} w_{l j} s_{j} \tag{3}
\end{equation*}
$$

It will be shown that the energy decreases when the state is changed according to the given rules, until a minimum is reached. It is required that the weight matrix is symmetric, which is achieved automatically by the Hebb learning rule for pattern recognition provided that the bounds -1 and 1 are used and the initial matrix is symmetric. The change in energy due to the evaluation of element $k$ is (all other elements remain unchanged)

$$
\begin{align*}
& \Delta E=E^{\prime}-E \\
& \text { \{using equatıon 1\} } \\
& =-\frac{1}{2} \sum_{\imath j} s_{\imath}^{\prime} w_{t j} s_{j}^{\prime}+\frac{1}{2} \sum_{\imath} s_{i} w_{\imath j} s_{\jmath} \\
& =-\frac{1}{2}\left(\sum_{i} s_{i}^{\prime} w_{i k} s_{k}^{\prime}+\sum_{j} s_{k}^{\prime} w_{k j} s_{j}^{\prime}+\sum_{i} s_{i} w_{i k} s_{k}+\sum_{j} s_{k} w_{k j} s_{j}\right) \\
& \left\{w_{l l}=0\right\} \\
& =-\frac{1}{2}\left(\sum_{i} s_{i} w_{i k} s_{k}^{\prime}+\sum_{j} s_{k}^{\prime} w_{k j} s_{j}+\sum_{i} s_{i} w_{i k} s_{k}+\sum_{j} s_{k} w_{k j} s_{j}\right) \\
& \text { \{collect terms\} } \\
& =-\frac{1}{2}\left(\sum_{i} s_{i} w_{t k}\left(s_{k}^{\prime}-s_{k}\right)+\sum_{j}\left(s_{k}^{\prime}-s_{k}\right) w_{k j} s_{j}\right) \quad \text { [rearranging] } \\
& =-\frac{1}{2}\left(s_{k}^{\prime}-s_{k}\right) \sum_{j}\left(w_{k j}+w_{j k}\right) s_{j} \\
& \left\{w_{u j}=w_{j u}\right\} \\
& =\left(s_{k}^{*} s_{k}\right) \sum_{J} w_{k j} s_{J} \\
& \text { \{using equation 3\} } \\
& =-\left(s_{k}^{\prime}-s_{k}\right) v_{k} \tag{4}
\end{align*}
$$

It is easy to see that $\Delta E=0$ when $s_{k}^{\prime}=s_{k}$. If the output of element $k$ is changed, than $s_{k}^{\prime}=-s_{k}$ and using equation 2 one finds $\Delta E=-2 s_{k}^{\prime} v_{k}<0$. Since $E$ is bounded, the system will evolve to a (local) minimum of $E$.

The observation that a Hopfield network searches a minimum energy in the system space leads to several new applications. The network is now used to find the best solution for some assignment problem instead of pattern recognition [Hopfield 85]. Well-known examples are the travelling salesman problem and job assignment for people with different rates for the tasks. An energy surface is defined that corresponds to the problem. An $\mathrm{N} \times \mathrm{N}$-network is used for a problem of size N . First, N N-flops are built. An N -flop is the N dimensional analogue of a flip-flop, it is a collection of N neurons which are only in a stable state when one of them is active. This is achieved by inhibitory connections between the neurons. These N -flops are used to achieve that each site is visited only once in the TS-problem or each task is assigned to exactly one person in the job assignment problem. The N -flops are interconnected with a strength depending on the problem; the distance between the cities or the ranking of people for the tasks. When the network is in a stable state the active neurons show the solution found by the network; the best path for the salesman or the optimum task assignment. Unfortunately, this optimum found can correspond to a local energy
minimum, then the assignment is a sub-optimal one. In some practical situations such a solution may be acceptable. Networks using continuous valued generator functions appear to avoid local minima better than the corresponding networks with binary generator functions. The derivation of the convergence shown above is not applicable for networks with continuous values.

## Application of a feedback network

We will illustrate the use of Hopfield networks by constructing a feedback net that searches a solution of the N -queens problem; place N queens on an $\mathrm{N} \times \mathrm{N}$ chess board, such that they cannot hit each other. Of course, an $\mathrm{N} \times \mathrm{N}$ network is used for this purpose. The connections are arranged such that the queens are repulsive; the potential of a field is decteased when it is covered by a queen on another field. Furthermore, there is a tendency in the generator function to place queens on the board when a field is not covered by any other queen. In order to try to avoid local minima, we use pseudo linear outputs and a diagonal term to decrease the changes between the various steps. In this way a hill climbing system is built. The asynchronous evolution function is used to compute each new state.

```
queen_generator :: generator
queen_generator v = clip (v+1/3)
queen_element :: element
queen_element weights = queen_generator.wsum weights
compute_state :: weights_net -> input_state -> net_state
compute_state = feed_back_asyn queen_element
compute_output :: weights_net -> input_state -> output_state
compute_output weights input = select_output (compute_state weights input)
initial_net :: [num] >> weights_net
initial_net [n] = [[[matrix_value i j n | i <- [0..(n^2-1)]] | j <- [0..(n^2-1)]]}
matrix_value position1 position2 dimension
= 1 , position1 = position2 Il diagonal term
=-1/2 ,i=k V j=IV |l rook-like attack
            abs (i-k) = abs (j-l) I| bishop-like attack
=0 , otherwise | no attack
    where (i,j) = coordinates position1 dimension
                        (k,l) = coordinates position2 dimension
```

```
coordinates field_number board_size
\(=\) (row,col)
    where row = field_number div board_size
    col = field_number mod board_size
```

```
high = 1.0
low = 0.0
nr_constant_neurons = 0
```

The network described can be used to change an initial configuration to a solution with the functions defined in appendix C. To show the behaviour of this network, the results of two evaluations are reproduced here. Note that the initial configurations are isomorphic upon a rotation over $90^{\circ}$. The increasing probability to place a queen on a field is indicated by the characters ' ', '.', ' + ', ' $*$ ', ' $q$ ' and ' Q '.


Unfortunately, the system gets frequently stuck in a local minimum: a stable state with too few queens on the board. The standard solution to this familiar problem is to add a noise term which enables the system to leave the local minima. This term is decreased each iteration and one hopes that the system cools down in the absolute energy minimum. The terminology is borrowed from physics due to the similarity with thermal noise in spin models.

## Avoiding local minima

The Hopfield network with binary values is isomorphic [Amit 85] to the Ising spin model for many coupled spin particles [Reif 65, Reichl 80]. There is a famous technique in physics to find the minimum energy state for such a system called simulated annealing. The system starts at a high temperature which implies that much noise is added when a new system state is computed; there is a large chance to change spin. While new states are computed, the temperature is slowly decreased. It can be shown that the expectation value of the system state is the minimum energy state. This technique is also used in feedback networks to leave local energy minima [Kirkpatrick 83]. Unfortunately, this does not imply that a feedback network will always find the minimum energy state when a slowly decreasing amount of noise is added during the iterations.

We will show that the addition of a decreasing noise term to the states helps the system to find solutions for the queens problem. A real simulated annealing
procedure requires binary neurons and a state change proportional to $\exp (-\Delta \mathrm{E} / \mathrm{T})$ for uphill state changes, state changes that reduce the energy are always accepted. The original Hopfield networks are the $T=0$ limit of this rule.

```
compute_state :: weights_net -> input_state -> net_state
compute_state = annealing 0.5 queen_element
annealing :: num -> element -> weights_net ->input_state->[output_state]
annealing t element [weights] input
= to_limit (disturb t (eval_one_by_one element weights) input)
disturb :: num -> (state->state) -> state -> [state]
disturb t eval state
= state : disturb (annealing_rate * t) eval new_state
    where new_state = eval[stateli + noise t (i+v)|i<-[0..#state-1]]
    v = sum state
```

For binary units it is better to flip to values in the state with a decreasing probability instead of adding the noise term.
noise :: num -> num -> num
noise t $s=t^{*}\left(1-\left(\right.\right.$ random $\left(23459^{*}\right.$ entier $\left.\left.\left.s\right) \bmod 101\right) / 50\right) \quad$ II pseudo random $[-t, t]$

## Avoiding local minima in the N -queens example

The result of the evaluation of second initial state, using this new network definition with an annealing rate of 0.8 , is shown below.


The network behaves better using this evaluation function. Although, it is most likely to find an absolute energy minimum, it cannot be assured that the system ends up in a state with N queens. We have not bothered ourselves with the optimization of the initial 'temperature' and annealing rate of this network. The current choice of parameters works for a number of cases, but certainly not always. The relatively poor performance of this network is not surprising; as indicated above local minima cannot be avoided by this network. It should only be used in situations where a sub-optimal solution is acceptable.

An effect similar to the addition of noise can be achieved by making large steps in the gradient decent algorithm. The generator function and the weights are changed compared to the first attempt.

| queen_generator $v=$ clip | $(v+3)$ | II a pseudo linear or, |
| :--- | :--- | :--- |
| queen_generator $v=$ sigmoid | $(v+3)$ | II a differentiable generator function |

matrix_value position1 position2 dimension

$$
\begin{array}{lll}
=-2 & , \text { position1 }=\text { position2 } & \text { || diagonal term } \\
=-2 & , i=k \vee j=I V \text { abs }(i-k)=a b s(j-1) \quad \text { || artack } \\
=0 & , \text { otherwise } & \text { || no artack } \\
\text { where } & \text { (i,j) }=\text { coordinates position1 dimension } \\
& (k, I)=\text { coordinates position2 dimension }
\end{array}
$$

Using the first queen_generator function listed above the evolution of the network becomes:


Using this network a correct solution is found for all start configurations isomorphic to the shown form. Andres reports very good results for these weights, but he uses the second generator function [Andres 90]. He shows a derivation for the form of the weights in the matrix from an appropriate energy function. The magnitude of these weights is set in a trial and error process. A network with a sigmoid generator function needs more iterations to come near to a good solution, but seems to avoid local minima a bit better. Wrong solutions caused by getting stuck in a local minimum appear to occur only in small networks [Andres, private communications]. The number of iterations needed appears to increase very slowly with the size of a network. The 10 -queens problem requires on average 69 iterations, the mean number of iterations needed to solve the 300 -queens is 205.

When this solution is compared with an ordinary backtrack solution of the 8 -queens problem in a functional language, the performance of the neural network is quite disappointing. The backtrack solution finds the 92 solutions to the 8 -queens problem within approximately the same time as it takes to compute the weight matrix for the neural network algorithm! The backtrack algorithm used is given below.

The search space for the backtrack algorithm is reduced by an efficient coding of occupied positions. A simple analysis shows that each row must contain exactly one queen. So, the solution can be described by a list of positions of the queens in each row. A further analysis shows that also each column must contain one queen. Hence, the solutions are permutations of the numbers [0..n-1], where bishop-like attacks are prevented. The empty list of queens is extended in all possible ways until all queens are placed.

| size | $==$ | num |
| :--- | :--- | :--- |
| solution | $==$ | II the size of the board |
| partial_solution | $==$ [position] | Il the position of the queen on each row |
| position | $==$ | num |$\quad$ II queen position on a row; column number

```
queens :: size -> [solution]
queens \(\mathrm{n}=\) extend [][0..n-1]
extend :: partial_solution \(\rightarrow\) [position] \(->\) [solution]
extend part unused
\(=\) [part \(\quad\), if unused \(=[]\)
= concat [ extend (q:part) (filter ( \(\sim q\) ) unused)
    \(\mid \mathrm{q}<-\) unused ; and [abs ( q - parti) ~= \(\mathrm{i}+1 \mid \mathrm{i}<-[0 . . \# p a r t-1]]]\), otherwise
    II try each remaining position in next row; test for bishop-like attacks
```


### 6.7 Discussion

Functional languages appear to be very well suited to describe neural networks, although the descriptions presented here are longer than the customary descriptions using constructive mathematics. Mathematics is also more convenient for the proofs shown in this chapter.

The implementation effort required for these network descriptions is small and it yields a number of advantages. The obtained specifications are executable, partial correct and complete. This in contradiction to the usal mathematical descriptions where, for instance, the bounds for the summations and the evaluation order of elements is not specified. Apart from these usual advantages there are some additional benefits; due to the used framework of description tools the similarities and differences between the described networks are made explicit. These tools appear to be suitable to describe a large number of networks as shown in this chapter. Furthermore, these tools are expected to be suited for the description of many other kinds of networks. The described neural networks have indeed many similarities; all networks discussed use McCulloch-Pitts neurons arranged in a layered structure. Each network uses a fixed type of elements. The only difference among the elements in a network is the weight vector used to specify the individual elements.

The executability of the neural networks specifications presented here, is particularly convenient to determine properties which cannot be deduced analytically, but must be determined empirically. The stochastic behaviour, found in many networks, can be simulated by a pseudo random generator. To obtain really stochastic behaviour, a random generator must be added to the functional language, although this spoils the referential transparency.

In spite of the very simple processing elements used and the limited possibilities to prove that the artificial neural networks will work, these networks appear to be able to perform a number of stiff tasks. But, the large number of teaching samples needed by perceptrons and the unpredictable changes in this number are a severe problem in the application of these neural network models in computer science. A single pass through the network requires much computation; usually quadratically increasing (or worse) with the size of the input pattern. It is in principle possible to construct special purpose hardware to speed up the evaluation of a neural network.

A LVQ network has more appealing properties; the learning speed is usually much higher and more constant. Moreover, the amount of computations required is smaller. Unfortunately it requires some initial knowledge about the number of classes required and the number of patterns in these classes; the number of vectors needed in a given class cannot be predicted. Moreover, the leaming algorithm used in these networks does not always work. It is even possible that a network which behaves totally correct 'forgets' what it has learned.

Since the artificial neural networks treated in this chapter can be modelled adequately in a functional programming language the set of functions that can be computed by these networks is smaller or equal to the functions expressible in a functional programming language. However, the character of programming is quite different. A neural network is trained to map similar inputs to similar outputs, a correct behaviour is only obtained by a very rich set of trainings samples. An ordinary program is a generally applicable algorithm.

## Chapter 7 Summary and Conclusion

The suitability of functional programming languages as a formalism for executable specifications is investigated in this thesis.

A specification is a detailed product description in a suited formalism. Specifications are used in computer science to describe the functional and operational behaviour of algorithms and machines. A good specification is consistent, correct and sufficiently complete. A suited specification formalism is unambiguous and yields clear and compact descriptions.

A prototype is an implementation that obeys (some aspects of) the specification. It is constructed in order to achieve confidence in of the suitability of the specified system and in the correctness of the specification.

Using a programming language as description formalism has a number of advantages. Programming languages have usually a well defined and clear semantics, yielding unambiguous and understandable descriptions. The implementation of the programming language can be used to check some correctness aspects of the specification. A compiler cannot examine whether a specification is correct, but it can at least spot syntax errors, undefined identifiers and type conflicts. Last but not least, the description is executable, so a specification is its own prototype.

Despite of these benefits programming languages are seldom used as description formalism. The expressive power of most programming languages appears to be too low to obtain a concise description. Due to the lack of suited description primitives much programming effort is needed. The size of the obtained description is too large to be clearly readable. Moreover, the distinction between the specification and the product vanishes.

Functional programming languages are advertised as high level languages with a great expressive power. Hence they seem to be suited as description formalism. We have tried the use of functional languages as description formalism in practice. On one hand we have compared descriptions in functional languages with some other description formalisms, on the other hand we have applied them in some complex real-life applications.

In chapter 2 a hierarchical machine description in the functional programming language Miranda is compared with a traditional description. It appears that the functional specification is a bit longer. This is caused by the layout rules we have imposed ourselves to obtain the clearest description. The functional description is more direct than the imperative description. Moreover, it is very understandable and appears to be a useful prototype.

This description method is applied with success in chapter 3 to a more complex concrete problem: the specification of an abstract imperative graph rewrite machine. This machine and its description are used in practice as an intermediate level in the implementation of functional languages. The obtained specification is actually the one and only definition of this machine. It is found that it is easy to experiment with slightly different instruction sets in this prototype.

In chapter 4 the suitability of functional languages for the description of program translations is investigated. A program translation in Miranda is compared with some conventional description methods. The specification in a functional language is at least as clear as the other descriptions.

The translation of the graph rewrite language Clean to code for the machine defined in chapter 3 is specified with such a specification in chapter 5. This specification was a useful guide-line for the construction of the compiler. Although not all optimizations implemented are described, it explains the kind of code generated very well. This description was used together with the prototype machine to execute Clean programs before the corresponding products were made.

The artificial neural network descriptions presented in chapter 6 are implementations of the neural models. The specification is a useful addition to the customary specification in a mathematical style. This mathematical specification gives some dependencies of values, but no constructive way to combine these dependencies to a reasonable efficient algorithm. Moreover, there are often inaccuracies, like omitted bounds for a summation, in these mathematical specifications. These mathematical specifications can be easily recognized as a part of the functional description. A set of description tools is presented and applied to describe a number of networks. The similarities and differences between the various networks are easily spotted in this uniform description.

In conclusion it can be stated that functional specifications are applied successfully in a number of real applications. Comparison with conventional descriptions shows that the overhead imposed by using a general purpose functional language instead of tailor made formalism is limited. The obtained specifications are clear descriptions and usable prototypes.

Little attention is given to formal proofs of the correctness of the specifications. Due to the size of the applications it is a tedious job to derive and verify such proofs. However, functional specifications are suited for the rewrite type of proofs shown in chapter 2 . The equivalency of terms can in fact be checked by the prototype implementation. The proofs in chapter 6 are given in ordinary
mathematics; there is still too little experience with symbolic manipulation of functional programs to perform these proofs with the functional specification.

The general purpose functional language Miranda is used for the specifications presented in this thesis. It appears to be reasonably suited as specification formalism. There are, however, two problems with this language. Due to the lack of rewrite semantics partially parametrized functions are not treated as proper root normal forms. This hampers the presentation of the instructions in our machine descriptions. Secondly, the Milner/Mycroft type system forces us to inject some superfluous constructors in the data structures used to represent programs for the transformations in the chapters 4 and 5.

# Appendix A Definition of the ABC-Machine 

## A. 1 The micro-instruction set

A. 2 The instruction set
A. 3 ABC-assembler

This appendix contains the complete specification of the ABC-machine and ABC-assembler. Moreover some additional functions to obtain a better usable prototype implementation are included. These topics are discussed in chapter 3.

## A. 1 The micro-instruction set

abstype astack, bstack, cstack, graphstore, descstore, pc, programstore, io with

## Definition of the A-stack

as_get $\quad::$ a_src -> astack -> nodeid
as_init $::$ astack
as_popn :: nr_args -> astack -> astack
as_push :: nodeid -> astack -> astack
as_pushn :: nodeid_seq -> astack -> astack
as_topn $::$ nr_args $->$ astack $->$ nodeid_seq
as_update :: a_dst -> nodeid -> astack -> astack

## Definition of the B-stack

| bs_copy | $::$ b_src $->$ bstack $->$ bstack |
| :--- | :--- |
| bs_get | $::$ b_src $->$ bstack $\rightarrow$ basic |
| bs_getB | $::$ b_src $>$ bstack $\rightarrow$ boolean |
| bs_getI | $:$ b_src $->$ bstack $\rightarrow$ int |
| bs_init | $::$ bstack |
| bs_popn | $:$ nr_args $->$ bstack $\rightarrow$ bstack |
| bs_push | $::$ basic $\rightarrow$ bstack $\rightarrow$ bstack |

```
bs_pushB :: boolean -> bstack -> bstack
bs_pushI :: int -> bstack -> bstack
bs_update :: b_dst -> basic -> bstack -> bstack
bs_addI :: bstack -> bstack
bs_decI :: bstack -> bstack
bs_eqB :: bstack > bstack
bs_eql :: bstack -> bstack
bs_eqBi :: boolean -> b_src -> bstack -> bstack
bs_eqIi :: int -> b_src -> bstack -> bstack
bs_gtl :: bstack -> bstack
bs_incl :: bstack -> bstack
bs_ltI :: bstack -> bstack
bs_muli :: bstack -> bstack
bs_subI :: bstack -> bstack
```


## Definition of the C-stack

```
\begin{tabular}{ll} 
cs_init & \(::\) cstack \\
cs_get & \(::\) c_src \(->\) cstack \(->\) instrid \\
cs_popn & \(::\) c_src \(->\) cstack \(->\) cstack \\
cs_push & \(\because:\) instrid \(->\) cstack \(\rightarrow\) cstack
\end{tabular}
```


## Definition of the Nodes

n_arg $\quad::$ node -> arg_nr -> arity -> nodeid
n_args $\quad::$ node $->$ arity $->$ nodeid_seq
n_arity $\quad::$ node $->$ arity
n_B :: node -> boolean
n_copy $\quad::$ node $->$ node $->$ node
n_descid $::$ node $\gg$ descid
n_entry $\quad::$ node $\rightarrow$ instrid
n_eq_arity $\quad::$ node $->$ arity $->$ boolean
n_eq_B $\quad::$ node $\rightarrow$ boolean $->$ boolean
n_eq_descid :: node $\rightarrow$ descid $->$ boolean
n_eq_I $::$ node $->$ int $>$ boolean
n_eq_symbol : : node $->$ node $->$ bool
n_fill $\quad:$ descid $->$ instrid $->$ args $->$ node $->$ node
n_fillB $\quad::$ descid $\rightarrow$ instrid $\rightarrow$ boolean $\rightarrow$ node $\rightarrow$ node
n_fillI $\quad::$ descid $->$ instrid $->$ int $->$ node $\rightarrow$ node
n_I $\quad:$ : node $->$ int
n_nargs $\quad::$ node $->$ nr_args $->$ arity $->$ nodeid_seq
n_setentry $::$ instrid $\rightarrow$ node $\rightarrow$ node

## Definition of the Graph Store

gs_get $\quad::$ nodeid $->$ graphstore $->$ node
gs_init :: graphstore
gs_newnode :: graphstore -> (graphstore, nodeid)
gs_update :: nodeid -> (node -> node) -> graphstore -> graphstore

## Definition of the Descriptors

d_ap_entry :: desc -> instrid
d_arity $\quad::$ desc -> num
d_name :: desc -> string

## Definition of the Descriptor Store

ds_get $\quad::$ descid $->$ descstore $->$ desc
ds_init $::$ [desc] -> descstore

## Definition of the Program Counter

pc_init $\quad:$ instrid
pc_next $\quad::$ instrid $\rightarrow$ instrid
pc_halt $\quad::$ instrid $\rightarrow$ instrid
pc_end :: instrid $->$ bool

## Definition of the Program Store

ps_get :: instrid -> programstore -> instruction
ps_init : : [instruction] -> programstore
Definition of the I/O-channels
io_init : : io
io_print $\quad::$ string $->$ io $->$ io

## Micro-instructions to enable a trace

```
show_as :: astack -> [char]
show_bs :: bstack >> [char]
show_cs :: cstack -> [char]
show_io :: i0 -> [char]
show_pc :: instrid -> [char]
show_gs :: graphstore -> descstore -> [char]
show_nds :: num -> [node] -> descstore -> [char]
show_nd :: node >> descstore -> [char]
```

The state of the ABC-machine

| state | $==$ (astack,bstack,cstack,graphstore,descstore,instrid,programstore,io) |
| :--- | :--- |
| instruction | $==$ state $\rightarrow$ state |
| ap_entry | $==$ instrid |
| args | $==$ nodeid_seq |
| arity | $==$ nat |
| boolean | $==$ bool |
| entry | $==$ instrid |
| int | $==$ num |
| name | $==$ [char] |
| nat | $==$ num |


| nodeid | $==$ | nat |
| :--- | :--- | :--- |
| nodeid_seq | $==$ | [nodeid] |
| label | $==$ | [char] |
| string | $==$ | [char] |
| a_dst | $==$ | nat |
| a_src | $==$ | nat |
| a_src1 | $==$ | nat |
| a_src2 | $==$ | nat |
| b_dst | $==$ | nat |
| b_src | $==$ | nat |
| c_src | $==$ | nat |
| nr_args | $==$ | nat |
| arg_nr | $==$ | nat |

## Implementation of the A-stack

astack $==$ [nodeid]
as_get target as
= asget 0 as
where asget $n(a: x)$
$=\mathbf{a} \quad$, if $n=$ target
$=$ asget $(n+1) x$, otherwise
asget n []
$=$ error ("Taking element "++show target++" of A-stack of length "++show n)
as_init =[]
as_popn n as
= aspopn n as
where aspopn 0 as $=$ as
aspopn $m$ (a:x) $=$ aspopn $(m-1) x$
aspopn m [ ]
= error ("Popping "++show n++" elements of A-stack of length "++show(\#as))
as_push nodeid as $=$ nodeid:as
as_pushn nodeids as = nodeids ++ as
as_topn 0 as = []
as_topn $n$ (x:r) $=x:($ as_topn ( $n-1$ ) r)
as_topn $n$ [] $=$ error "as_topn: taking too much elements"
as_update 0 nodeid (a:x) = nodeid: $x$
as_update $n$ nodeid ( $a: x$ ) = a : as_update ( $n-1$ ) nodeid $x$
as_update n nodeid as = error "A-stack update error"

## Implementation of the B-stack

bstack == [basic]
bs_get target bs
$=b s g e t 0 \mathrm{bs}$
where bsget $n(b: x)$

$$
\begin{aligned}
& =\mathrm{b} \\
& =\text { bsget }(n+1) \times, \text { if } n=\text { target } \\
& \text { bsget } n[] \\
& =\text { error ("Taking element "++show target++" of B-stack of length "++show } n \text { ) }
\end{aligned}
$$

bs_getB $n$ bs $=b$ where Bool $b=b s \_$get $n$ bs
bs_getInbs = $i$ where Inti =bs_getnbs
bs_copy $n$ bs $=\left(b s \_g e t n b s\right):$ bs
bs_init =[]
bs_popn $n$ bs
= bspopn n bs
where bspopn 0 bs = bs bspopn $m(b: x)=$ bspopn $(m-1) x$ bspopn m [ ]
= error ("Popping "++show n++" elements of B-stack of length "++show(\# bs))
bs_push basic bs = basic:bs
bs_pushB b bs $\quad=($ Bool b):bs
bs_pushIibs = (lnti):bs
bs_update 0 basic (b:x) = basic:x
bs_update $n$ basic ( $\mathrm{b}: \mathrm{x}$ ) $=\mathrm{b}$ : bs_update ( $\mathrm{n}-1$ ) basic x
bs_update $n$ basic bs $=$ error " $B$-stack update error"
bs_addI ((Int i1):(Int i2):r) = ( $\operatorname{lnt}(11+i 2)): r$
bs_decI ((Int i1):r) $=(\operatorname{lnt}(i 1-1)): r$
bs_eqB ((Bool b1):(Bool b2):r) = (Bool (b1=b2)):r
bs_eqBi b $\quad$ bs $\quad=$ bs_push (Bool (b=(bs_getB $\boldsymbol{n}$ bs))) bs
bs_eqI ((Int i1):(Int i2):r) $=$ (Bool ( $\mathrm{i} 1=\mathrm{i} 2)$ ): :
bs_eqli inbs = bs_push (Bool (i=(bs_getl $n$ bs))) bs
bs_gtI ((Int i1):(Int i2):r) = (Bool (i1>i2)):r
bs_incl ((Inti1):r) $=(\operatorname{lnt}(i 1+1)): r$
bs_ItI ((Int i1):(Int i2):r) = (Bool (i1<i2)):r
bs_mulI ((Inti1):(Int i2):r) = (Int (i1*i2)):r
bs_subI ((Int i1):(Int i2):r) = (Int (i1-i2)):r

## Implementation of the C-stack

```
cstack == [Instrid]
Cs_init = []
Cs_get target cs
= csget O cs
    where csget n (c:x)
                            =c , if n= target
                            = csget (n+1) x . otherwise
```

```
            csget n[]
            = error ("Taking element "++show target++" of C-stack of length "++show n)
cs_popn n cs
    = cspopn n cs
    where cspopn Ocs = cs
                cspopn m (c:x) = cspopn (m-1)x
                aspopn m [ ]
            = error("Popping "++show n++" elements of C-stack of length "++show(#cs))
CS_push c CS = c:CS
Implementation of the Nodes
node \(\left.::=\begin{array}{ll}\text { Node descid instrid args } \\ & \text { Basic descid instrid basic } \\ & \text { Empty }\end{array}\right) \mid\)
basic : := Int int
Bool boolean
n_arg node n arity
\(=\) args ! ( \(\mathrm{n}-1\) ) if arity \(>=n\)
= error ("taking arg " ++ show n ++ " from " ++ show node), otherwise where args = n_args node arity
n_args (Node descid entry args) arity
\(=\) args \(\quad\), if arity \(=\) \# args
= error ("wrong arity in n_args: " ++ show arity ++ " (Node " ++ show descid++" "++show entry++" "++show args++")"), otherwise
n_args node arity
= error ("n_args: No args in node: " ++ show node)
\(n_{\text {_arity }}\) (Basic descid entry basic) \(=0\)
n_arity (Node descid entry args) = \# args
n_arity Empty \(\quad=\) error "Arity of an empty node is not defined"
n_B (Basic descid entry (Bool b)) = b
n_B node = error ("N_B: No boolean in node:" ++ show node)
n_copy new old = new
n_descid (Node descid entry args) = descid
n_descid (Basic descid entry basic) = descid
n_descid Empty \(\quad=\) error "No descid in an Empty node!"
n_entry (Node descid entry args) = entry
n_entry (Basic descid entry basic) \(=\) entry
n_entry Empty \(\quad=\) error "No entry in an Empty node!"
\begin{tabular}{ll} 
n_eq_arity & node \(n\) \\
n_eq_B & node \(b\) \\
n_eq_descid node descid & \(=n_{\_}\)arity node \(=\boldsymbol{n}\) node \(=b\) \\
& \(=n_{1}\) descid node \(=\) descid
\end{tabular}
```

| n_eq_symbol (Node descid1 entry1 args1) (Node descid2 entry2 args2) |  |
| :---: | :---: |
| n_eq_symbol (Basic descid1 entry1 basic1) (Basic descid2 entry2 basic2) |  |
|  | $=$ descid1 $=$ descid2 \& basic1 = basic2 |
| n_eq_symbol node1 node2 | $=$ False |
| $n$ fill desc entry args node | = Node desc entry args |
| n_fillB desc entry b node | = Basic desc entry (Bool b) |
| n_filli desc entry i node | = Basic desc entry (Int i) |
| n_I (Basic descid entry (Int i)) | $=\mathbf{i}$ |
| n_I node | = error ("n_I: No integer in node: " ++ show node) |
| n_nargs node arg_count arity | = take arg_count (n_args node arity) |
| n_setentry newentry (Node descid entry args) = Node descid newentry args |  |
| n_setentry newentry (Basic descid entry basic) = Basic descid newentry basic |  |
| n_setentry newentry Empty | = error "Cannot set entry of Empty node!" |

## Implementation of the Descriptors

desc $\quad:=$ Desc ap_entry arity name
d_ap_entry (Desc ap_entry arity name) = ap_entry
d_arity (Desc ap_entry arity name) = arity
d_name (Desc ap_entry arity name) = name
Implementation of the Descriptor Store

```
descid == num
descstore == [desc]
```

ds_get target ds
$=\operatorname{dsget} 0 \mathrm{ds}$
where dsget $\mathrm{n}(\mathrm{d}: \mathrm{x})$
$=\mathrm{d} \quad$, if $\mathrm{n}=$ target
$=\operatorname{dsget}(n+1) x$, otherwise
dsget $n$ []
$=$ error ("Taking descriptor "++show target++" of store of size "++show n)
ds_init descriptors = descriptors

Implementation of the Graph Store

```
graphstore == ([node], nat)
gs_get nodeid (nds,free)
    = gsget (nodeid-free-1) nds
        where gsget 0 (n:ns) = n
                            gsget m (n:ns) = gsget (m-1) ns
                                gsget m[] = error ("unexisting nodeid : " ++ show nodeid)
gs_init = ([ ], store_size)
```

```
store_size nat
store_size = 100
Il some natural number indicatng the size
gs_newnode (nds,0) = error "graph-store is full"
gs_newnode (nds,free) = ((Empty nds,tree-1),free)
gs_update noderd f (nodes,free)
= (update place nodes f,free) .ff place<=store_size free
= error ("gs update error node "++show nodeid++" not existent"),otherwise
    where place = noderd - free -1
update :: nat -> [node] -> (node->node) -> [node]
update 0 (node nodes) f = f node nodes
update n (node nodes) f = node (update ( }n-1\mathrm{ ) nodes f)
Implementation of the Program Store
programstore == [location]
location = l instruction
```



```
ps_get target ps
= psget 0 ps
    where psget locus (l instr rest)
                                    = instr , , locus = target
                                    = psget (locus+1) rest , otherwise
                                    psget locus [ ] = error ("Program counter outside program " ++ show target)
Implementation of the Program Counter
instrid \(==\) num
pc == instrid
pc_end instrid =instrid < 0
pc_init =0
pc_halt instrid = -1
pc_next instrid = instrid + 1
Implementation of the I/O channels
\begin{tabular}{ll}
10 & \(==\) [char] \\
\(10 \_\)init & \(=[]\)
\end{tabular}
10_print string output = output ++ string
symbol_to_string node -> desc -> string
symbol_to_string (Basic descid ap_entry (Int 1)) desc = show I
symbol_to_string (Basic descid ap_entry (Bool b)) desc = showb
symbol_to_string (Node descid entry args)(Desc ap_entry arity name) = name
```

```
show_state : : state -> [char]
show_state (as,bs,cs,gs,ds,pc,ps,io)
\(=\) "output : " ++ show_io io ++ " \(n^{n}\) ++
    "pc : " ++ show_pc pc ++ " \(n\) " ++
    "A-stack: " ++ show_as as ++ " \(n^{\prime \prime}++\)
    "B-stack: " ++ show_bs bs ++ " \(n\) " ++
    "C-stack: " ++ show_cs cs \(+{ }^{1 / n "++}\)
    "Graph : (nodeid : node): \(\quad\) " \(n\) " ++ show_gs gs ds
```

| show_as as | $=$ show as |  |
| :--- | :--- | :--- |
| show_bs bs | $=$ show bs |  |
| show_cs cs | $=$ show cs |  |
| show_io io | $=$ io |  |
| show_gs (nds,free) ds | $=$ show_nds (free+1) nds ds |  |
| show_pc pc | $=$ show pc | , if pc $>=0$ |

show_nds $i$ (n:nds) ds = show $i+{ }^{\prime \prime}$ : "++show_nd n ds++"n"++show_nds (i+1) nds ds
show_nds i[]ds = []
show_nd (Node descid entry args) ds
= d_name (ds_get descid ds)++" "++show entry++" "++show args
show_nd (Basic descid entry basic) ds = show entry ++"n++ show basic
show_nd Empty ds = "Empty"
cond : : bool -> * -> * -> *
cond $b$ then else
$=$ then, if $b$
= else, otherwise

## A. 2 The instruction set

Type definitions of the instructions

| add_args | :: a_src -> nr_args -> a_dst -> instruction |
| :---: | :---: |
| create | : : instruction |
| del_args | : : a_src -> nr_args $->$ a_dst $->$ instruction |
| dump | :: string -> instruction |
| eq_desc | : : descid -> a_src -> instruction |
| eq_desc_arity | $::$ descid $->$ arity $->$ a_src $->$ instruction |
| eq_symbol | : : a_src1 $\rightarrow$ a_src2 $->$ instruction |
| eqB | : : instruction |
| eqB_a | : : bool $\rightarrow$ a_src $->$ instruction |
| eqB_b | : : bool -> b_src -> instruction |
| eqI | : : instruction |
| eqI_a | : : int $->$ a_src $->$ instruction |
| eqI_b | : : int $->$ b_src $->$ instruction |
| fill | $::$ descid $->$ nr_args $\rightarrow$ instrid $->$ a_dst $->$ instruction |
| fill_a | : : a_src -> a_dst $->$ instruction |
| fill ${ }^{\text {l }}$ | :: bool $->$ a_dst $->$ instruction |


| fillB_b | : : b_src -> a_dst -> instruction |
| :---: | :---: |
| fill | $::$ int $\quad \rightarrow$ a_dst $->$ instruction |
| fillı ${ }^{\text {b }}$ | $::$ b_src $->$ a_dst $->$ instruction |
| get_desc_arity | :: a_src -> instruction |
| get_node_arity | : : a_src -> instruction |
| halt | $:$ : instruction |
| jmp | $::$ instrid -> instruction |
| jmp_eval | $::$ instruction |
| jmp_false | : : instrid -> instruction |
| jmp_true | : $:$ instrid -> instruction |
| jsr | : : instrid -> instruction |
| jsr_eval | :: instruction |
| no_op | : : instruction |
| pop_a | : : nr_args -> instruction |
| pop_b | : : nr_args $->$ instruction |
| print | : : string $->$ instruction |
| print_symbol | $::$ a_src -> instruction |
| push_a | :: a_src -> instruction |
| push_ap_entry | $::$ a_src $->$ instruction |
| push_arg | :: a_src -> arity $->$ arg_nr $->$ instruction |
| push_arg_b | : : a_src $\rightarrow$ - instruction |
| push_args | : : a_src -> arity $->$ nr_args $->$ instruction |
| push_args_b | :: a_src -> instruction |
| push_b | : $:$ b_src -> instruction |
| pushB | : : bool -> instruction |
| pushB_a | :: a_src -> instruction |
| pushi | : : int -> instruction |
| pushi_a | :: a_src -> instruction |
| repl_args | : $:$ arity $\rightarrow>$ nr_args $->$ instruction |
| repl_args_b | : : instruction |
| th | : : instruction |
| set_entry | : : instrid -> a_dst -> instruction |
| update_a | $::$ a_src $\rightarrow$ a_dst $\rightarrow$ instruction |
| update_b | : : b_src $\rightarrow$ b ${ }^{\text {b dst }}$-> instruction |
| addI | : : instruction |
| decI | : $:$ instruction |
| gtI | : : instruction |
| incl | : : instruction |
| ItI | : : instruction |
| muli | :: instruction |
| subI | : : instruction |

## Implementation of the instructions

```
add_args a_src nr_args a_dst (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs',ds,pc,ps,io)
    where as' = as_popn nr_args as
        gs' = gs_update dstid (n_fill descid entry newargs) gs
        dstid = as_get a_dst as
        srcid = as_get a_src as
        node = gs_get srcid gs
        descid = n_descid node
        entry = n_entry node
        arity = n_arity node
        newargs = n_args node arity ++ as_topn nr_args as
```

create (as,bs,cs,gs,ds,pc,ps,io)
$=(a s ', b s, c s, g s ', d s, p c, p s, i o)$
where as' = as_push nodeid as
(gs',nodeid) = gs_newnode gs
del_args a_src nr_args a_dst (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs',ds,pc,ps,io)
where as' = as_pushn newargs as
gs' $\quad=$ gs_update dstid (n_fill descid entry newargs) gs
dstid $=$ as_get a_dst as
srcid $=$ as_get a_src as
node $\quad=$ gs_get srcid gs
descid $=$ n_descid node
entry $=$ n_entry node
newargs = n_nargs node (arity-nr_args) arity
arity $=$ n_arity node
dump string (as,bs,cs,gs,ds,pc,ps,io)
$=$ ( $\mathrm{as}, \mathrm{bs}, \mathrm{cs}, \mathrm{gs}, \mathrm{ds}, \mathrm{pc}, \mathrm{ps}, \mathrm{io}$ )
where io' $=0 \_$print (" $n^{n}++$ string $+{ }^{m}{ }^{n} n^{\prime}++$ state) io
state $\quad=$ show_state (as,bs,cs,gs,ds,pc,ps,io)
eq_desc descid a_src (as,bs,cs,gs,ds,pc,ps,io)
$=(\mathrm{as}, \mathrm{bs}, \mathrm{cs}, \mathrm{gs}, \mathrm{ds}, \mathrm{pc}, \mathrm{ps}, \mathrm{io})$
where bs' $=$ bs_pushB equal bs
equal $=$ n_eq_descid node descid
node $\quad=$ gs_get nodeid gs
nodeid $=$ as_get a_src as
eq_desc_arity descid arity a_src (as,bs,cs,gs,ds,pc,ps,io)
$=(\mathrm{as}, \mathrm{bs}, \mathrm{cs}, \mathrm{gs}, \mathrm{ds}, \mathrm{pc}, \mathrm{ps}, \mathrm{io})$
where bs' = bs_pushB equal bs
equal $=$ (n_eq_descid node descid) \& (n_eq_arity node arity)
node $\quad=$ gs_get nodeid gs
nodeid $=$ as_get a_src as

```
eq_symbol a_src1 a_src2 (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_pushB equal bs
        equal = n_eq_symbol node1 node2
        node1 = gs_get nodeid1 gs
        node2 = gs_get nodeid2 gs
        nodeid1 = as_get a_src1 as
        nodeid2 = as_get a_src2 as
eqB (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_eqB bs
eqB_a bool a_src (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_pushB equal bs
        equal = n_eq_B (gs_get nodeid gs) bool
        nodeid = as_get a_src as
eqB_b bool b_src (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_eqBi bool b_src bs
eqI (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_eqI bs
eqI_a int a_src (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_pushB equal bs
        equal = n_eq_l (gs_get nodeid gs) int
        nodeid = as_get a_src as
eqI_b int b_src (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_eqli int b_src bs
fill desc nr_args entry a_dst (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs',ds,pc,ps,io)
    where as' = as_popn nr_args as
        gs' = gs_update nodeid (n_fill desc entry args) gs
        nodeid = as_get a_dst as
        args = as_topn nr_args as
fill_a a_src a_dst (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs,gs',ds,pc,ps,io)
    where gs' = gs_update nodeid_dst (n_copy node_src) gs
        node_src = gs_get nodeid_src gs
        nodeid_dst = as_get a_dst as
        nodeid_src = as_get a_src as
```

fillB bool a_dst (as,bs,cs,gs,ds,pc,ps,10)

```
= (as,bs,cs,gs',ds,pc,ps,o)
    where gs' = gs_update noderd (n_fillB bool_desc rnf_entry bool) gs
        nodeld = as_get a_dst as
flllB_b b_src a_dst (as,bs,cs,gs,ds,pc,ps,ıo)
= (as,bs,cs,gs',ds,pc,ps,וo)
    where gs' = gs_update noderd (n_fillB bool_desc rnf_entry bool)gs
        bool = bs_getB b_src bs
        noderd = as get a_dst as
```

filli int a_dst (as,bs,cs,gs,ds,pc,ps,io)
$=(\mathrm{as}, \mathrm{bs}, \mathrm{cs}, \mathrm{gs}, \mathrm{ds}, \mathrm{pc}, \mathrm{ps}, 10)$
where $g s^{\prime} \quad=\quad$ gs_update nodeid (n_filli int_desc rnf_entry int) gs
nodeıd = as_get a_dst as
filli_b b_src a_dst (as,bs,cs,gs,ds,pc,ps,o)
= (as,bs,cs,gs',ds,pc,ps,io)
where gs' = gs_update noderd (n_filli int_desc rnf_entry int) gs
int = bs_getl b_src bs
nodeid = as_get a_dst as
get_desc_arity a_src (as,bs,cs,gs,ds,pc,ps,10)
$=(\mathrm{as}, \mathrm{bs}, \mathrm{cs}, \mathrm{gs}, \mathrm{ds}, \mathrm{pc}, \mathrm{ps}, 10)$
where bs' = bs_pushI arity bs
arity $=$ d_anty (ds_get descid ds)
descid $=$ n_descid (gs_get nodeıd gs)
nodeid $=$ as_get a_src as
get_node_arity a_src (as,bs,cs,gs,ds,pc,ps,io)
$=\left(a s, b s{ }^{\prime}, c s, g s, d s, p c, p s, 10\right)$
where bs' = bs_pushi arity bs
arity $=$ n_arity (gs_get noderd gs)
noderd $=$ as_get a_src as
halt (as,bs,cs,gs,ds,pc,ps,io)
$=(\mathrm{as}, \mathrm{bs}, \mathrm{cs}, \mathrm{gs}, \mathrm{ds}, \mathrm{pc}, \mathrm{ps}, \mathrm{וo})$
where $p c^{\prime} \quad=p c$ halt $p c$
رmp address (as,bs,cs,gs,ds,pc,ps,ıo)
= (as,bs,cs,gs,ds,pc',ps,10)
where $\mathrm{pc}^{\prime}=$ address
Jmp_eval (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs,gs,ds,pc',ps,10)
where $p c^{\prime} \quad=\quad$ n_entry ( $g s$ _get nodeld $g s$ )
nodeıd = as_get 0 as

```
jmp_false address (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc',ps,io)
    where pc' = cond (~bool) address pc
        bool = bs_getB 0 bs
        bs' = bs_popn 1 bs
jmp_true address (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs',cs,gs,ds,pc',ps,io)
    where pc' = cond bool address pc
        bool = bs_getB 0 bs
        bs' = bs_popn 1 bs
jsr address (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs',gs,ds,pc',ps,io)
    where pc' = address
        cs' = cs_push pcecs
jsr_eval (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs,cs',gs,ds,pc',ps,io)
    where pc' = n_entry (gs_get nodeid gs)
        nodeid = as_get 0 as
        cs' = cs_push pccs
no_op (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs,cs,gs,ds,pc,ps,io)
pop_a n (as,bs,cs,gs,ds,pc,ps,io)
    = (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_popn n as
pop_b n (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_popn n bs
print string (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs,cs,gs,ds,pc,ps,io')
    where io' = io_print string io
print_symbol a_src (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs,cs,gs,ds,pc,ps,io')
    where io' = io_print string io
        nodeid = as_get a_src as
        node = gs_get nodeid gs
        string = symbol_to_string node desc
        desc = ds_get (n_descid node) ds
push_a a_src (as,bs,cs,gs,ds,pc,ps,io)
    = (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_push nodeid as
        nodeid = as_get a_src as
```

```
push_ap_entry a_src (as,bs,cs,gs,ds.pc,ps,וo)
= (as,bs,cs',gs,ds,pc,ps,io)
    where cs' = cs_push (d_ap_entry (ds_get descid ds)) cs
        descid = n_descid (gs_get noderd gs)
        nodeıd = as_get a_src as
push_arg a_src arity arg_nr (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_push arg as
        arg = n_arg (gs_get nodeid gs) arg_nr arity
        noderd = as_get a_src as
push_arg_b a_src (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_push arg as
        arg = n_arg (gs_get nodeid gs) arg_nr arity
        nodeid = as_get a_src as
        arg_nr = bs_getI O bs
        arity = bs_getI 1 bs
push_args a_src arity nr_args (as,bs,cs,gs,ds,pc,ps,o)
    = (as',bs,cs,gs,ds,pc,ps,10)
    where as' = as_pushn args as
        args = n_nargs (gs_get nodeld gs) nr_args arity
        nodeid = as_get a_src as
push_args_b a_src (as,bs,cs,gs,ds,pc,ps,10)
= (as',bs,cs,gs,ds,pc,ps,וo)
    where as' = as_pushn args as
        args = n_nargs (gs_get noderd gs) nargs arity
        nargs = bs_getI O bs
        noderd = as_get a_src as
        arity = bs_getl 1 bs
push_b b_src (as,bs,cs,gs,ds,pc,ps,ı0)
    = (as,bs',cs,gs,ds,pc,ps,10)
    where bs' = bs_push basic bs
        basic = bs_get b_src bs
pushB bool (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_pushB bool bs
pushB_a a_src (as,bs,cs,gs,ds,pc,ps,o)
    = (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_pushB bool bs
        bool = n_B (gs_get nodeid gs)
        noderd = as_get a_src as
pushI int (as,bs,cs,gs,ds,pc,ps,ıo)
    = (as,bs',cs,gs,ds,pc,ps,ı0)
    where bs' = bs_pushI int bs
```

```
pushI_a a_src (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_pushI int bs
    int = n_l (gs_get nodeid gs)
    nodeid = as_get a_src as
repl_args arity nr_args (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_pushn args (as_popn 1 as)
        args = n_nargs (gs_get nodeid gs) nr_args arity
        nodeid = as_get 0 as
repl_args_b (as,bs,cs,gs,ds,pc,ps,io)
= (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_pushn args (as_popn 1 as)
    args = n_nargs (gs_get nodeid gs) nr_args arity
        nodeid = as_get 0 as
        arity = bs_getI 0 bs
        nr_args = bs_getI 1 bs
rtn (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs',gs,ds,pc',ps,io)
    where pc' = cs_get 0 cs
        cs' = cs_popn 1 cs
set_entry entry a_dst (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs,cs,gs',ds,pc,ps,io)
    where gs' = gs_update nodeid (n_setentry entry) gs
        nodeid = as_get a_dst as
update_a a_src a_dst (as,bs,cs,gs,ds,pc,ps,io)
    = (as',bs,cs,gs,ds,pc,ps,io)
    where as' = as_update a_dst nodeid as
        nodeid = as_get a_src as
update_b b_src b_dst (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_update b_dst basic bs
        basic = bs_get b_src bs
addI (as,bs,cs,gs,ds,pc,ps,io)
= (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_addI bs
decl (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_decI bs
gtI (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs',cs,gs,ds,pc,ps,io)
    where bs' = bs_gtI bs
```

$$
\begin{aligned}
& \text { incI (as,bs,cs,gs,ds,pc,ps,io) } \\
& =\left(a s, b s^{\prime}, c s, g s, d s, p c, p s, i o\right) \\
& \text { where bs' } \quad=\text { bs_incI bs }
\end{aligned}
$$

It (as,bs,cs,gs,ds,pc,ps,io)
$=(a s, b s ', c s, g s, d s, p c, p s, i o)$
where bs' = bs_挰bs
mull (as,bs,cs,gs,ds,pc,ps,io)
$=(a s, b s ', c s, g s, d s, p c, p s, i o)$
where bs' = bs_mull bs
subI (as,bs,cs,gs,ds,pc,ps,io)
$=$ (as,bs',cs,gs,ds,pc,ps,io)
where bs' = bs_subl bs

## The driver

```
boot : : ([instruction],[desc]) -> state
boot (program,descriptors)
= (as,bs,cs,gs,ds,pc,ps,io)
    where \(p c=p c_{\text {_init }}\)
        as = as_init
        bs = bs_init
        CS = CS_init
        gs = gs_init
        ps = ps_init program
        io \(=\) lo_init
        ds = ds_init descriptors
    fetch_cycle : : state -> state
    fetch_cycle (as,bs,cs,gs,ds,pc,ps,io)
    = (as,bs,cs,gs,ds,pc,ps,io), pc_end pc
    \(=\) fetch_cycle (currinstr (as,bs,cs,gs,ds,pc',ps,io)), otherwise
        where \(p c^{\prime}=\) pc_next \(p c\)
            currinstr = ps_get pc ps
step_cycle (as,bs,cs,gs,ds,pc,ps,io)
    = show_state (as,bs,cs,gs,ds,pc,ps,io), pc_end pc
    = show_state (as,bs,cs,gs,ds,pc,ps,io) ++
    step_cycle (currinstr (as,bs,cs,gs,ds,pc',ps,io')), otherwise
    where pc' = pc_next pc
        currinstr = ps_get pc ps
        io' \(=\) lo_init
int_desc \(=0\)
bool_desc = 1
rnf_entry = 1
```


## A. 3 ABC-assembler

ABC -assembler is represented by an AST. The mapping from ABC-assembler to ABC -instructions is a simple abstract program transformation.


| Print | string |
| :---: | :---: |
| Print_symbol | a_src |
| Push_a | a_src |
| Push_ap_entry | a_src |
| Push_arg | a_src arity arg_nr |
| Push_arg_b | a_sic |
| Push_args | a_src arity arg_nr |
| Push_args_b | a_src |
| Push_b | nat |
| PushB | bool |
| PushB_a | a_src |
| Pushi | int |
| Pushi_a | a_src |
| Repl_args | arity $7 r_{-}$args |
| Repl_args_b |  |
| Rtor |  |
| Set_entry | label a_dst |
| Update_a | a_src a_dst |
| Update_b | b_src b_dst |
| AddI |  |
| DecI |  |
| GII |  |
| Incl |  |
| LtI |  |
| Muli |  |
| Subi |  |

assemble : : assembler $->$ ([instruction],[desc])
assemble statements
= (translate statements loc_counter sym_table,desc_table statements sym_table)
where loc_counter $=0$
desc_counter $=0$
sym_table $=$ collect statements loc_counter desc_counter
collect : : assembler -> nat -> nat -> sym_table
collect (Label l:r) Ic dc = (I,lc,Lab_sym):collect r lc dc
collect (Descriptor desc_label red_label arity name:r) Ic dc
$=$ (desc_label,dc,Desc_sym):collect r ic (dc+1)
collect (instr:r) Ic dc $=$ collect $r(l c+1) d c$
collect [] lc dc $=$ []
look_up : : label -> sym_type -> sym_table -> nat
look_up lab $t$ ((name,n,sym_type):r)
$=n \quad$ if lab = name $\& \mathbf{t}=$ sym_type
= look_up labtr , otherwise
look_up lab t [ ] = error ("In\nlabel" ++ lab ++ " not defined as " ++ show t ++ " n " )
sym_table $==$ [(name,nat,sym_type)]
sym_type ::= Lab_sym | Desc_sym
desc_table: :assembler -> sym_table -> [desc]
desc_table (Descriptor desc_label ap_entry arity name:r) sym_table
= (Desc ap_addr arity name):desc_table r sym_table
where ap_addr = look_up ap_entry Lab_sym sym_table
desc_table (instr:r) sym_table = desc_table $r$ sym_table
desc_table [] sym_table = []
translate : :assembler -> nat -> sym_table -> [instruction]

translate ( $\mathrm{Br}_{\mathbf{r}}$ false $\boldsymbol{n}$
$=$ jmp_false ( $\mathrm{I}+\mathrm{n}+\mathrm{n}+1$ )
translate ( $\mathrm{Br}_{\mathbf{\prime}}$ true n
= jmp_true ( $\mathrm{Ic}+\mathrm{n}+1$ )
translate (Dump string
= dump string
translate (Add_args a_src n a_dst
= add_args a_src $n$ a_dst
translate (Create
= create
translate (Del_args a_src n a_dst
= del_args a_src $n$ a_dst
translate (Eq_desc desc_label a_src
= eq_desc desc_addr a_src

| translate | r) ic |
| :---: | :---: |
|  | r) lc |
| translate | $r$ lc |
|  | r) lc |
| : translate | $r$ ( $\mathrm{lc}+$ |
|  | r) Ic |
| : translate | r (lc+ |
|  | r) lc |
| : translate | $r$ ( $\mathrm{lc}+$ |
|  | r) lc |

: translate r ( $\mathrm{Ic}+1$ )
$: \quad$ r) lc
:translate r ( $\mathrm{lc}+1$ )
: r) lc
:translate r (lc+1)
$: \quad$ r) ic
:translate r ( $\mathrm{Ic}+1$ )
r) Ic
:translate r (lc+1)
sym_table
sym_table
sym_table
sym_table
sym_table
sym_table
sym_table
sym_table
sym_table
sym_table
sym_table
sym_table sym_table sym_table sym_table sym_table sym_table sym_table sym_table sym_table where desc_addr = look_up desc_label Desc_sym sym_table translate (Eq_desc_arity desc_label arity a_src :
r) Ic sym_table
$=$ eq_desc_arity desc_addr arity a_src : translate r (Ic+1) sym_table where desc_addr = look_up desc_label Desc_sym sym_table
translate (EqB : r) ic sym_table
= eqB
:translate r (lc+1)
sym_table
translate (EqB_a bool a_src
= eqB_a bool a_src
translate (EqB_b bool b_src
= eqB_b bool b_src
translate (EqI
= eqI
translate (EqI_a int a_src
= eqI_a int a_src
translate (EqI_b int b_src
= eql_b int b_src
translate (Fill desc_lab nr_args entry_lab a_dst :
= fill desc_addr nr_args entry_addr a_dst
: r) lc
:translate $\mathbf{r}$ ( $\mathrm{lc}+1$ )
r) lc
:translate r ( $\mathrm{lc}+1$ )
$: \quad$ r) lc
:transiate r ( $\mathrm{lc}+1$ )
$: \quad$ r) lc
: translate r (lc+1)
$: \quad$ r) lc :translate r (lc+1) :translate r ( $\mathrm{lc}+1$ ) : translate r ( $\mathrm{lc}+1$ )
Desc sym sym_table where desc_addr = look_up desc_lab Desc_sym sym_table entry_addr = look_up entry_lab Lab_sym sym_table
sym_table sym_table sym_table sym_table sym_table sym_table sym_table sym_table sym_table sym_table sym_table sym_table
translate (Fill_a a_src a_dst
= fill_a a_src a_dst
translate (FillB bool a_dst
= fillB bool a_dst
translate (FillB_b b_src a_dst
= fillB_b b_src a_dst
translate (Filli int a_dst
= fillI int a_dst
translate (Filli_b b_src a_dst
= filli_b b_src a_dst
translate (Get_desc_arity a_src
= get_desc_arity a_src
translate (Get_node_arity a_src
= get_node_arity a_src
translate (Halt
= halt
translate (Jmp label
= jmp address
where address = look_up label Lab_sym sym_table
translate (Jmp_eval : r) ic sym_table
= jmp_eval
translate (Jmp_false label
= jmp_false address where address = look_up label Lab_sym sym_table
translate (Jmp_true label
= jmp_true address
where address = look_up label Lab_sym sym_table:
$\begin{array}{lllll}\text { translate (Jsr label } & : & \text { r) } & \text { lc } & \text { sym_table } \\ =\text { jsr address } & \text { : translate } & \text { r } & \text { (lc+1) } & \text { sym_table }\end{array}$
translate (Jsr_eval
= jsr_eval
translate (No_op
= no_op
translate (Pop_a nat
= pop_a nat
translate (Pop_b nat
= pop_b nat
translate (Print string
= print string
translate (Print_symbol a_src
= print_symbol a_src
translate (Push_a a_src
= push_a a_src
translate (Push_ap_entry a_src
= push_ap_entry a_src
translate (Push_arg a_src arity arg_nr
= push_arg a_src arity arg_nr
: r) lc sym_table
: translate r ( $\mathrm{Ic}+1$ ) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) ic sym_table
: translate r ( $\mathrm{Ic}+1$ ) sym_table
: r) lc sym_table
: translate $r$ ( $\mathrm{lc}+1$ ) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: translate r (lc+1)
sym_table
sym_table
sym_table $\begin{array}{llll}: & \text { r) } & \text { lc } & \text { sym_table } \\ \text { : translate } & \text { r } & \text { (lc+1) } & \text { sym_table }\end{array}$ $\begin{array}{llll}: & \text { r) } & \text { lc } & \text { sym_table } \\ \text { : translate } & \text { r } & \text { ( } \mathrm{lc}+1) & \text { sym_table }\end{array}$ : r) lc sym_table : translate r (lc+1) sym_table
: r) lc sym_table
: translate $\mathbf{r}(\mathrm{lc}+1)$
sym_table
: r) lc sym_table
: translate r (Ic+1) sym_table
: r) lc sym_table
: translate $r$ ( $\mathrm{I}+\mathrm{c}+1$ )
sym_table
: r) lc sym_table
: translate r (ic+1) sym_table
: r) lc sym_table
: translate r (Ic+1) sym_table
: r) lc sym_table
: translate r (Ic+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) ic sym_table
: translate $\mathbf{r}$ ( $\mathrm{Ic}+1$ )
translate (Push_arg_b a_src
= push_arg_b a_src
translate (Push_args a_src arity nr_args
= push_args a_src arity nr_args
translate (Push_args_b a_src
= push_args_b a_src
translate (Push_b nat
= push_b nat
translate (PushB bool
= pushB bool
translate (PushB_a a_src
= pushB_a a_src
translate (PushI int
= pushI int
translate (PushI_a a_src
= pushI_a a_src
translate (Repl_args arity nr_args
= repl_args arity nr_args
translate (Repl_args_b
= repl_args_b
translate (Rtn
= m
translate (Set_entry label a_dst
= set_entry address a_dst
where address = look_up label Lab_sym sym_table
translate (Update_a a_sre a_dst
= update_a a_src a_dst
translate (Update_b a_src a_dst
= update_b a_src a_dst
translate (AddI
= addI
translate (DecI
= decI
translate (GtI
= gtI
translate (IncI
= incl
translate (LtI
$=1 \mathrm{t} 1$
translate (MulI
= muli
translate (SubI
= subl
translate ( $x$
= error ("\n\nUnknown assembler statemant: " ++ show x )
translate [ ] lc sym_table = [ ]

[^7] -
: r) lc sym_table
: translate $\mathbf{r}$ ( $\mid c+1$ ) sym_table
: r) lc sym_table
: translate $\mathbf{r}$ ( $\mathrm{lc}+1$ ) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
$: \quad$ r) lc sym_table
: translate $\mathbf{r}$ ( $\mathrm{lc}+1$ ) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) lc sym_table
: translate $\mathbf{r}$ ( $\mid \mathbf{c}+1$ ) sym_table
: r) ic sym_table
: translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) Ic sym_table
: translate $r$ ( $\mathrm{Ic}+1$ ) sym_table
: r) ic sym_table : translate r (lc+1) sym_table
$: \quad$ r) lc sym_table : translate r (lc+1) sym_table
: r) lc sym_table : translate r (lc+1) sym_table
: - r)
:translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) Ic sym_table
: translate r (lc+1) sym_table
: r) lc sym_table : translate r (lc+1) sym_table
: r) lc sym_table : translate r (Ic+1) sym_table
: r) lc sym_table : translate r (lc+1) sym_table
: r) lc sym_table
: translate r (lc+1) sym_table
: r) lc sym_table : translate r (lc+1) sym_table
r) lc sym_table
listing : : assembler -> [char]
listing statements
= list statements loc_counter sym_table
where loc_counter $=0$ desc_counter $=0$
sym_table = collect statements loc_counter desc_counter
list : : assembler -> nat -> sym_table -> [char]
list (Label name r) ic sym_table
$=$ "nitt"++name ++ ":" ++ list r lc sym_table
list (Descriptor lab ap_entry arity name
r) Ic
sym_table
= "|n|t"++lab++"("++desc_num++"):ItDescriptor "++ap_entry++
" ("++ ap_addr++") "++show arity++" "++name ++ list r lc sym_table where ap_addr = show (look_up ap_entry Lab_sym sym_table) desc_num = show (look_up lab Desc_sym sym_table)
list $(\mathrm{Br} \mathrm{n}$
r) Ic sym_table

list (Br_false n : r) Ic sym_table
$=$ In Ic ++ "||Br_false " ++ show n ++ "|nittitmp_false "
++ show $(\mathrm{l} \mathrm{c}+\mathrm{n}+1) \quad+\quad$ list $r$ ( $\mathrm{l} \mathrm{c}+1)$ sym_table
list (Br_true n
$=$ In Ic ++ "||Br_true " ++ show $n++$ "niltitJmp_true " ++ show (Ic+n+1)

+ list $r$ ( $\mathrm{lc}+1$ ) sym_table
list (Eq_desc desc_label a_src
= In Ic++"Eq_desc "++desc_label++" ("++ show desc_addr++") "++show a_src + list r (Ic+1) sym_table where desc_addr = ook_up desc_label Desc_sym sym_table
list (Eq_desc_arity desc_label anty a_src
: r) Ic
sym_table
$=$ In IC++"Eq_desc_arity "++desc_label++" ("++desc_addr $++{ }^{4}$ ) "++show arity++" "++show a_src + list r (Ic+1) sym_table where desc_addr = show (look up desc_label Desc_sym sym_table)
list (Fill desc_lab nr_args entry_lab a_dst :
r) Ic
sym table
$=$ In Ic++"FIll "++desc_lab++" ("++desc_addr++") "++ show nr_args++
" "++entry_lab++" ("++ entry_-addr++") "++show a_dst ++ list $r$ ( $\mathrm{lc}+1$ 1) sym_table
where desc_addr = show (look_up desc_lab Desc_sym sym_table) entry_addr = show (look_up entry_lab Lab_sym sym_table)
list (Jmp label
r) Ic
= In Ic ++ "Jmp " ++ label ++ " (" ++ show address ++")" + list r (Ic+1)
sym_table
where address = look_up label Lab_sym sym_table
list (Jmp_false label
r) ic sym_table
$=$ In Ic ++ "Jmp_false "++label++" ("++show address ++ ")"++ list rell (ic+1) sym_table where address = look_up label Lab_sym sym_table
list (Jmp_true label
r) Ic sym_table
$=$ In Ic++"Jmp_true "++label++" ("++ show address ++ ")"++ list r (Ic+1) sym_table where address $=$ look_up label Lab_sym sym_table
list (Jsr label
$=\ln \mathrm{IC}++$ "Jsr " ++ label ++ " $($ " ++ show address ++ ")" $++\quad$ list $\quad$ r ( $\mathrm{Ic}+1)$ sym_table where address = look_up label Lab_sym sym_table
list (Set_entry label a_dst
r) ic
sym_table
$=$ In Ic++"Set_entry "++label++
" ("++show address++")"++show a_dst $\quad+$ list $r(i c+1)$ sym_table where address = look_up label Lab_sym sym_table list (instr $: \quad$ r) ic sym_table
$=$ In ic ++ show instr list [ ] Ic sym_table
= []
In :: num -> [char]
In $\mathrm{lc}=$ " $n \mathrm{n}$ " + show $\mathrm{lc}++$ mitt"


# Appendix B Translating Clean to ABC-code 

## B. 1 Reduction entries

B. 2 Code for the left handside of a rule alternative
B. 3 Code for the right handside of a rule alternative
B. 4 Auxiliary functions

A complete and executable specification of the translations of Clean to ABCcode is given in this appendix. The translation is explained and illustrated with a number of examples in chapter 5 . The Clean program is represented by the AST defined in section 5.3. The generated ABC-assembler is stored in the AST defined in appendix A. The machine specification in given in that chapter can be used to execute the translated Clean programs.

```
bindings == node_id -> offset || Associates an index in the A-stack frame to a node-id
offset == num |l An index in the A-stack frame
asp == num
state == node_id -> status || Associates the status of the node to a node-id
status || The states possible
    := Undefined | |l This node-id is not defined in the current scope
    Created | |l The node is created but is still empty
    Unknown | | The node is filled; it is unknown if the graph is a rnf
    InRnf |l This node is in root normal form
compiler:: clean -> abc_assembler
compile :: rewrite_rule -> abc_assembler
compile (TypedRule typealts alts)
= fun_descriptor Ihs ++
    prepare_args lhs ++
    alt_entries alts 1 ++
    gen_type_error symbol_id (# args) (# alts)
    where lhs = get_lhs (hd typealts)
```

NODE annots nodeid (Function symbol_id) args = hd lhs
compile (ConsRule typealts alts)
$=$ fun_descriptor Ihs ++
prepare_args ths ++
alt_entries alts $1 \quad+$
fill_with_rnf symbol_id (\# args) (\# alts)
where Ihs = get_lins (hd typealts)
NODE annots nodeid (Function symbol_id) args = hd ins
compile (UntypedRule alts)
= fun_descriptor Ihs ++
prepare_args lhs ++
alt_entries alts $1 \quad+$
fill_with_rnf symbol_id (\# args) (\# alts)
where lins = get_lhs (hd alts)
NODE annots nodeid (Function symbol_id) args = hd lhs
compile (TypeRule typealts)
= descriptors typealts
fun_descriptor :: graph -> abc_assembler
fun_descriptor (NODE annots nodeid (Function symbol_id) args:r)
= [ Descriptor
symbol_id
(apply_label symbol_id)
(\# args)
symbol_id
descriptors :: rule -> abc_assembler
descriptors (alt:rest)
$=$ cons_descriptor rhs ++
fill_code rhs ++
descriptors rest
where rhs = get_rhs alt
descriptors []
= []
cons_descriptor :: graph -> abc_assembler
cons_descriptor (NODE annots nodeid (Constructor symbol_id) args:r)
= [ Descriptor
symbol_id
apply_lab
arity
symbol_id ]
where apply_lab
$=$ rnf_lab $\quad$, if arity $=0$
= apply_label symbol_id , otherwise
arity $=$ \# args
cons_descriptor (NODE annots nodeid (Function symbol_id) args:r)
= []

```
fill_code graph -> abc_assembler
fill_code (NODE annots nodeid (Constructor symbol_id) args r)
= [ Label (apply_label symbol_ıd)
    Fill symbol_id arity rnf_lab arity
    Pts
= []
    where arity = \# args
fill_code (NODE annots nodeid (Function symbol_Id) args r)
= []
```

prepare_args graph -> abc_assembler
prepare_args (NODE annots noderd (Function symbol_ıd) args alts)
= [ Label (node_label symbol_ıd)
Set_entry cycle_lab 0
pushargs 0 arity
[ Label (apply_label symbol_ıd) ] +
reduce_strict_args args 0
where arity = \# args
reduce_strict_args arg's -> a_src -> abc_assembler
reduce_strict_args (arg args) n
$=$ reduce_arg $n++$ reduce_strict_args args $(n+1) \quad$ if is_strict arg
$=\quad$ reduce_strict_args args $(n+1)$, otherwise
reduce_strict_args [] $\quad n=$ []
reduce_arg a_src -> abc_assembler
reduce_arg a_src

|  | Jsr_eval |  |  | , if $\mathbf{a}$ _src $=0$ |
| :---: | :---: | :---: | :---: | :---: |
| = | Push_a | a_src |  |  |
|  | Jsr_eval |  |  |  |
|  | Pop_a | 1 |  | otherwise |

fill_with_rnf symbolid -> arity -> num -> abc_assembler
fill_with_rnf symbol_id arity nr_alts
$=$ [] , if arity $=0$
$=[$ Label (alt_label symbol_Id (nr_alts+1))
Fill symbol_ıd arity rnf_lab arity
R
gen_type_error symbolid -> arity -> num -> abc_assembler
gen_type_error symbol_ıd arity $\mathrm{nr}_{-}$alts
= []
, if arity $=0$
$=[$ Label (alt_label symbol_ıd (nr_alts +1 ))
Jmp type_error
, otherwise
, otherwise

## B. 1 Reduction entries

alt_entries rule $\rightarrow$ num $->$ abc_assembler
alt_entries (alt r) alt_number = alt_entry alt alt_number++alt_entries $r$ (alt_number +1 )
alt_entries [ ] alt_number = [ ]
alt_entry :: rulealt -> num -> abc_assembler
alt_entry (Rewrite lhs rhs) alt_number
= match_code ++ contractum rhs bindings states asp
where (match_code,bindings,states,asp) = match ths alt_number
alt_entry (Redirect lhs annots nodeid) alt_number
= match_code ++ redirection annots nodeid bindings states asp
where (match_code,bindings,states,asp) = match lhs alt_number

## B. 2 Code for the left handside of a rule alternative

match :: graph -> num -> (abc_assembler,bindings,states,asp)
match [NODE annots nodeid (Function symbol_id) args] alt_number
= (code, bindings', states', asp')
where code $=$ [Label label ] + match_code bindings = bind_args args asp empty_bindings states = record_args args initial_status asp $=$ arity arity = \# args label $=$ alt_label symbol_id alt_number next_alt = alt_label symbol_id (alt_number +1 ) (match_code,bindings',states',asp')
= match_args args bindings states asp next_alt asp arity
match_args $::$ arg's $->$ bindings $->$ states $->$ asp $->$ label $->$ offset $->$ arity $->$ (abc_assembler,bindings,states,asp)
match_args (Nodeld an nodeid:rest) bindings states asp next offset arity = match_args rest bindings states asp next (offset-1) arity
match_args (Term ans (NODE n_ans id sym args):rest) binds states asp next offset arity = (code, bindings", states", asp")
where code = bring_in_rnf annots a_src + + match_code ++ match_rest a_src = asp-offset (match_rest,bindings",states",asp")
= match_args rest bindings' states' asp' next (offset-1) arity (match_code,bindings',states',asp')
= match_arg (NODE (ans ++ n_ans) id sym args)
binds (record id InRnf states) asp next a_src arity
match_args [ ] bindings states asp next_alt offset arity
= ([ ],bindings,states,asp)
match_arg :: tree -> bindings -> states -> asp -> label -> a_src -> arity ->
(abc_assembler,bindings,states,asp)
match_arg (NODE anns id (INTval i) [ ]) binds states asp next a_src arity
$=$ (code ,binds, states, asp)
where code
= typecheck
[ EqI_a ia_src ] ++
escape_false (asp-arity) next
typecheck
= []
= [ Eq_desc int_symbolid a_src ] ++ escape_false (asp-arity) next , otherwise
match_arg (NODE ans id (BOOLval b) [ ]) binds states asp next a_src arity
$=($ code ,binds, states, asp)
where code
\(\left.=\begin{array}{ll}typecheck <br>

EqB_a ba_src\end{array}\right]\)| ++ |
| :--- |
| ++ |

typecheck
= []
if member anns is BOOL
$=[$ Eq_desc bool_symbolid a_src $]+$ escape_false (asp-arity) next ,otherwise
match_arg (NODE annots id symbol args) binds states asp next a_src arity
= (code, bindings", states", asp")
where code

$$
\begin{array}{rlr}
= & {[\text { Eq_desc_arity (symbol_id symbol) n_arity a_src] }} & ++ \\
& \text { escape_false (asp-arity) next } \\
& \text { pushargs a_src n_arity } \\
& \text { match_sub_args } \\
\text { bindings' }^{\prime}= & \text { bind_args args asp' binds } \\
\text { states' }^{\prime}= & \text { record_args args states } \\
\text { asp }^{\prime}= & \text { asp + n_arity } \\
\text { n_arity }= & \text { \# args } \\
\text { (match_sub_args,bindings",states",asp") } \\
= & \text { match_args args bindings' states' asp' next asp' arity }
\end{array}
$$

asp' $\quad=$ asp + n_arity
n_arity = \# args

## B. 3 Code for the right handside of a rule alternative

contractum :: graph -> bindings -> states -> asp -> abc_assembler
contractum (top: node_defs) bindings states asp
= build_shared_nodes ++ root top bindings' states' asp'
where (build_shared_nodes,bindings',states',asp')
= shared_nodes node_dets (bind top_id 0 bindings) states asp
NODE annots top_id symbol args $=$ top
shared_nodes :: [tree] -> bindings -> states ->asp->(abc_assembler,bindings,states,asp)
shared_nodes node_defs bindings states asp
= (code, bindings", states", asp")
where code = create_cycle_nodes ++ fill_nodes_code
(create_cycle_nodes,bindings',states',asp')
= cycle_nodes node_defs bindings states asp
(fill_nodes_code,bindings",states",asp")
= fill_nodes node_dets bindings' states' asp'
cycle_nodes :: [tree] -> bindings -> states -> asp -> (abc_assembler,bindings,states,asp) cycle_nodes [] bindings states asp
= ([ ],bindings,states,asp)
cycle_nodes (NODE annots nodeid sym args:rest) bindings states asp
= (code, bindings", states", asp") . if member annots OnACycle
= cycle_nodes rest bindings states asp ,otherwise
where code $=$ [ Create + code_rest
bindings' = bind nodeid asp' bindings
states' $=$ record nodeid Created states
asp' = asp + 1 (code_rest,bindings",states",asp") = cycle_nodes rest bindings' states' asp'
fill_nodes :: [tree] -> bindings -> states -> asp -> (abc_assembler,bindings,states,asp)
fill_nodes [ ] bindings states asp
$=([]$,bindings,states,asp)
fill_nodes (node: rest) bindings states asp
= (code, bindings", states", asp")
where code = creation + fill_code + fill_rest_code creation
= []
$=[$ Create
bindings'

$$
=\text { bindings }
$$

$$
=\text { bind nodeid asp' bindings , otherwise }
$$ asp'

$$
=\text { asp } \quad \text {. if previously_defined }
$$

$$
=\text { asp }+1 \quad \text {, otherwise }
$$

a_src = asp' - bindings' nodeid
(fill_code,states') = fill_node node bindings' states asp' a_src
(fill_rest_code,bindings",states",asp") = fill_nodes rest bindings' states' asp'

$$
\text { NODE annots nodeid sym args } \quad=\text { node }
$$

$$
\text { previously_defined } \quad=\text { states nodeid }=\text { Created }
$$

fill_node :: tree -> bindings -> states -> asp -> a_src -> (abc_assembler,states)
fill_node (NODE annots id (INTval i) [ ]) bindings states asp a_src
$=$ (code, states')
where code $=$ [ Filli ia_src ] states' = record id InRinf states
fill_node (NODE annots id (BOOLval b) [ ]) bindings states asp a_src
$=$ (code, states')
where code $=$ [ FillB ba_src states' = record id InAnf states
fill_node (NODE ans nid (Function sid) args) binds state asp asrc
$=$ fill_strict (NODE ans nid (Function sid) args) binds state asp asrc, member ans Strict
= (code, states") , otherwise
where code = build_args_code ++ [ Fill sid arity (node_label sid) (asrc+arity) ] (build_args_code,states") = build_args args binds states' asp
states' = record nid Unknown state
arity = \# args
fill_node (NODE annots nid (Constructor sid) args) binds states asp asrc
= (code, states")
where code = build_args_code ++ [ Fill sid arity rnf_lab (asrc + arity) (build_args_code,states") = build_args args binds states' asp states' $=$ record nid InRnf states arity = \# args
fill_strict :: tree -> bindings -> states -> asp -> a_src -> (abc_assembler,states)
fill_strict (NODE anns nodeid (Function sid) args) bindings states asp 0
= (code, new_states)
where code = build_args_code ++ [ Jsr (reduction_label sid)
new_states $=$ record nodeid InRnf states'
(build_args_code,states') = build_args args bindings states asp
fill_strict (NODE annot id (Function sid) args) bindings states asp asrc
$=$ (code, states)
where code

| $=\left[\begin{array}{l}\text { Push_a asrc } \\ \text { fill_code } \\ {[\text { Pop_a }} \\ \text { (fill_code,states) }\end{array}\right.$ | $]$+ |
| :--- | :---: |

$=$ fill_strict (NODE annot id (Function sid) args) bindings states (asp+1) 0
fill_strict node bindings states asp a_src
= fill_node node bindings states asp a_src
build_args :: arg's -> bindings -> states -> asp ->(abc_assembler,states)
build_args [ ] bindings states asp
$=$ ([ ],states)
build_args (arg:args) bindings states asp
= (code, states")
where code = args_code ++ arg_code
(args_code,states') = build_args args bindings states asp
(arg_code,states") = build_arg arg bindings states' (asp+\#args)
build_arg :: arg -> bindings $->$ states $->$ asp $->$ (abc_assembler,states)
build_arg (Nodeld annots nodeid) bindings states asp
$=$ (push ++ evaluation , new_states ) , if reduction_needed
= (push , states ) ,otherwise
where new_states $=$ record nodeid InRnf states
push $\quad$ [ Push_a (asp-bindings nodeid) ]
evaluation $=[$ Jsr_eval ]
reduction_needed = member annots Strict \& ~(states nodeid=InRnf)
build_arg (Term annots node) binds states asp
$=$ (creation ++ eval_code , strict_states ) , if member annots Strict
$=($ creation ++ fill_code , fill_states ) , otherwise
where creation $=$ [ Create ]
(eval_code,strict_states) $=$ fill_strict node binds states $(a s p+1) 0$
(fill_code,fill_states) = fill_node node binds states (asp+1) 0

```
root \(:\) : tree -> bindings \(->\) states \(->\) asp \(->\) abc_assembler
root (NODE annots id (INTval i) [ ]) bindings states asp
= pop_args asp ++
    [ Fill! io
    Rto
root (NODE annots id (BOOLval b) [ ]) bindings states asp
= pop_args asp ++
    [ FillB bo
    Pto
root (NODE annots nid (Constructor sid) args) bindings states asp
= build_args_code ++
    [ Fill sid arity rnf_lab (asp+arity) ] ++
    pop_args asp ++
    [ Ritn ]
    where (build_args_code,states') = build_args args bindings states asp
        arity = \#args
root (NODE annots nid (Function sid) args) bindings states asp
= build_args_code +
    clean_up (\# args) asp +
    [ Jmp (reduction_label sid) ]
    where (build_args_code,states') = build_args args bindings states asp
redirection \(::\) annots \(\rightarrow\) node_id \(->\) bindings \(\rightarrow\) states \(\rightarrow\) asp \(\rightarrow\) abc_assembler
redirection annots nodeid bindings states asp
\(=[\) Fill_a newroot asp
    Pop_a asp
    Rtn
\(=\) update_astack newroot (asp-1) +
    pop_args (asp-1) +
    [ Jsr_eval
        Fill_a 01
        Pop_a 1
        Rit
    where newroot = asp-(bindings nodeid)
```


## B. 4 Auxiliary functions

```
is_strict : : arg -> bool
is_strict (Nodeld annots nodeid) = member annots Strict
is_strict (Term annots node ) = member annots Strict
get_lhs : : rulealt -> graph
get_Ins (Rewrite lis rhs) = lhs
get_liss (Redirect lhs annots nodeid) = lhs
get_rhs : : rulealt -> graph
get_rhs (Rewrite Ihs rhs) = rhs
```

```
get_args : : tree -> arg's
get_args (NODE annots nodeid symbol args) = args
bring_in_mf : : annots -> a_src -> abc_assembler
bring_in_rnf annots a_src
= [] , if member annots Strict
= reduce_arg a_src , otherwise
escape_false :: nr_args -> label -> abc_assembler
escape_false nr_args label
\(=\) [ Jmp_false label ] , if nr_args =0
\(=\left[\begin{array}{ll}\text { Br_true } 2\end{array}\right.\)
    Pop_a nr_args
    Jmp label
pop_args : : nr_args -> abc_assembler
pop_args nr_args
= []
= [ Pop_a nr_args
, if nr_args = 0
] , otherwise
update_astack : : a_src -> a_dst -> abc_assembler
update_astack a_src a_dst
= []
= [ Update_a a_src a_dst
if a_src = a_dst
] , otherwise
pushargs : : a_src -> arity -> nr_args -> abc_assembler
pushargs a_src arity nr_args
= []
= [ Push_args a_src arity nr_args ]
```

if nr_args = 0
, otherwise

```
bind_args : : arg's -> offset -> bindings -> bindings
bind_args (Nodeld annots nodeid:rest) \(n\) old
\(=\) bind_args rest ( \(n-1\) ) (bind nodeid \(n\) old)
bind_args (Term annots (NODE annots' nodeid sym args):rest) n old
\(=\) bind_args rest ( \(n-1\) ) old
, if nodeid = ""
\(=\) bind_args rest ( \(n-1\) ) (bind nodeid \(n\) old) , otherwise
bind_args [ ] \(n\) bindings
= bindings
bind : : node_id \(->\) offset \(->\) bindings \(->\) bindings
bind nodeid offset bindings name
= offset , if name = nodeid
= bindings name , otherwise
empty_bindings : : bindings
empty_bindings nodeid = error ("Nodeld not bound in this stack frame: "++nodeid)
record_args : : arg's \(->\) states \(\rightarrow\) states
record_args (Nodeld annots id:rest) states
= record_args rest (record id InRnf states) , if member annots Strict
= record_args rest (record id Unknown states) , otherwise
```

```
record_args (Term annots (NODE n_annots id sym args):rest) states
    = record_args rest (record id InRnf states) , if member (annots++n_annots) Strict
    = record_args rest (record id Unknown states) , otherwise
record_args [] states
    = states
record: : node_id -> status \(->\) states \(->\) states
record nodeid status states name
    = states name \(\quad\) if nodeid \(={ }^{n n} \quad\) Il no node-id
    = status , if name = nodeid
    = states name , otherwise
initial_status : : states
initial_status name \(=\) Undefined
clean_up : : num -> num -> abc_assembler
clean_up keep 0 = []
clean_up 0 drop \(=[\) Pop_a drop ]
clean_up keep drop = [ Update_a (keep-1) (keep+drop-1)] ++ clean_up (keep-1) drop
node_label : : symbolid -> label
node_label symbolid = "n_" ++ symbolid
apply_label : : symbolid -> label
apply_label symbolid = "a_" ++ symbolid
alt_label : : symbolid \(->\) num \(->\) label
alt_label symbolid \(n=\) symbolid ++ show \(n\)
reduction_label : : symbolid -> label
reduction_label symbolid = alt_label symbolid 1
cycle_lab: : label
cycle_lab = "_cycle"
rnf_lab: : label
rnf_lab = "_rnf"
type_error : : label
type_error = "type_error"
symbol_id : : symbol -> symbolid
symbol_id (Symbol id) = id
symbol_id (Function id) \(=\) id
symbol_id (Constructor id) = id
symbol_id symbol \(=\) error ("No id in the symbol: " ++ show symbol)
int_symbolid : : symbolid
int_symbolid = "INT"
bool_symbolid : : symbolid
bool_symbolid = "BOOL"
```


# Appendix C Artificial Neural Networks 

## C. 1 Network description tools

C. 2 A feedback network to solve the N -queens problem

This appendix contains a complete definition of a set of tools to describe and investigate the behaviour artificial neural networks. These tools are used to describe a number of network models in chapter 6 . Only the specification of the network for the N -queens problem is not covered completely in the text, so the complete definition is shown here.

## C. 1 Network description tools

| weights_net | $=$ [weights_layer] | Il the synaps weights in a network |
| :---: | :---: | :---: |
| weights_layer | $=$ [weights] | Il the synaps weights in a layer |
| weights | $=$ = [weight] | $\\|$ the weights defining one synapse |
| weight | num |  |
| input_state | [ [input_value] | I\| list of simultaneous input_values |
| input_value | value |  |
| output_state | = [output_value] | II list of simultaneous output_values |
| output_value | value |  |
| net_state | $=$ = [state] | II the states of all layers in a net |
| state | [value] | Il the states of neurons in a layer |
| target_state | [target_value] | II the target output_state |
| target_value | value |  |
| potential | num | II the internal potential of a neuron |
| value | num | II used binary or continuously by application |
| vector | = [num] |  |
| vectors | $=$ [vector] |  |
| vector_id | = num |  |
| class | = [num] | If list to be compatible with output_state |
| target_class | $=$ = class | II used in learning vector quantisation |


| synapse | $==$ | weights $\rightarrow$ input_state $->$ potential |
| :--- | :--- | :--- |
| generator | $==$ | potential $\rightarrow$ output_value |
| element | $==$ | weights $\rightarrow$ neuron |
| neuron | $==$ | input_state $\rightarrow$ output_value |

## Functions determining the evaluation order

evaluate the state of of a layered network; state! 0 is the input state
feed_forward :: element $->$ weights_net $\rightarrow$ input_state $->$ net_state
feed_forward element weights input
= state
where state
= (input ++ constant_neurons):
[eval_layer element (weights!!) (state!!) ++constant_neurons | $\mathrm{I}<-$ [0..\#weights -1]]
constant_neurons :: state
constant_neurons = rep nr_constant_neurons high
feed_back_syn :: element -> weights_net -> input_state -> [output_state]
feed_back_syn element weights input
= to_limit(iterate(select_output.feed_forward element weights) input)
feed_back_asyn :: element $->$ weights_net $->$ input_state $->$ [output_state]
feed_back_asyn element [weights] input
= to_limit (iterate (eval_one_by_one element weights) input)
annealing :: num $->$ element $->$ weights_net ->input_state->[output_state]
annealing t element [weights] input
= to_limit (disturb t (eval_one_by_one element weights) input)
disturb :: num -> (state->state) $->$ state $->$ [state]
disturb t eval state
= state : disturb (annealing_rate*t) eval new_state
where $v \quad=$ sum state
$\Pi \quad=\#$ state
new_state $=$ eval $\left[\right.$ state $!i+$ noise $\left.t\left(n^{*} i+v\right) \mid i<-[0 . . n-1]\right]$ annealing_rate $=0.8$
noise :: num -> num $->$ num
noise ts $=t^{*}\left(1-\left(\right.\right.$ random $\left(23459^{*}\right.$ entier $\left.\left.\left.s\right) \bmod 101\right) / 50\right) \|$ pseudo random within $[-t, t]$
eval_layer $::$ element $->$ weights_layer $->$ input_state $->$ output_state
eval_layer element layer input = [element weights input | weights <- layer]
select_output :: net_state -> output_state
select_output state
= output_state , if nr_constant_neurons $=0$
= take (\#output_state-nr_constant_neurons) output_state , otherwise where output_state $=$ last state
eval_one_by_one :: element -> weights_layer -> input_state -> output_state
eval_one_by_one element layer input = eval_neurons element input layer [0..\# layer -1]
Il evaluate the neurons in a feed back network one after another in pseudo random order
eval_neurons :: element -> state -> weights_layer -> [num] -> state
eval_neurons element state w_layer [] = state ++ constant_neurons
eval_neurons element state w_layer indices
= eval_neurons element new_state w_layer rest
where $n=$ random (1993+entier (sum state)) mod \# indices
index $=$ indices ! $n$
rest $\quad=$ remove $n$ indices
new_state $=$ update index new_value state
new_value = element (w_layer!index) state
II Update one neuron at each iteration; taken pseudo random of the list of indices
random :: num -> num
random seed $=\left(16807^{*}\right.$ seed $) \bmod 2147483647 \quad| | 2147483647=2 \wedge 31-1$
learn :: weights_net $->$ [input_state] $->$ [target_state] $\rightarrow$ weights_net
learn weights (sample:samples) (target:targets)
= learn (adapt_weights weights sample target) samples targets
learn weights [][]
$=$ weights

## Functlons determining the specific network

| high | :: value | II The upper limit of the values |
| :---: | :---: | :---: |
| low | :: value | II The lower limit of the values |
| nr_constant_neurons | :: num | II The number of constant neurons |
| compute_state | :: weights_net $->$ input_state | -> [state] |
| compute_output | :: weights_net $->$ input_state | -> output_state |
| adapt_weights | :: weights_net -> input_state | -> target_state -> weights_net |
| gain | :: num | I\| Determines the learing rate |
| a | :: num | II Determines the size of the initial weights |
| initial_net | :: [num] -> weights_net |  |
| show_elem | :: value -> char |  |
| show_net | : weights_net -> [char] |  |

## The synaptlc functions

```
wsum :: synapse
wsum = inner_product
inner_product :: synapse
inner_product v1 v2 = sum (map2 (*) v1 v2) |l or sum [v1li* v2!i|i<-[0..# v1 -1]]
neocognitron_synapse :: synapse
neocognitron_synapse weights input
= (1+excitation)/(1+inhibition) - 1
    where excitation = sum[weights!i*inputi| i<-[0..#weights-1];weightsli>0]
        inhibition = -sum[weights!i*input!]|<-[0..#weights-1];weights!i<0]
```

The generator functions

$$
\begin{aligned}
& \text { hard_limit :: generator } \\
& \text { hard_limit } v \\
& =\text { high, if } v>\text { threshold } \\
& =\text { low ,otherwise }
\end{aligned}
$$

sgn :: generator
$\operatorname{sgn} v$
$=1$, if $v>0$
$=-1$. otherwise
sigmoid :: generator
sigmoid $\mathrm{v}=$ (high-low) $/\left(1+\exp \left(\right.\right.$-sharpness $\left.\left.{ }^{*} v\right)\right)+$ low
sigmoid' :: generator II the derivate of sigmoid
sigmoid' $v=\left((\text { high-low })^{*}\right.$ sharpness*exp (-sharpness* $\left.\left.{ }^{*}\right)\right) /\left((1+\exp (- \text { sharpness* } v))^{\wedge} 2\right)$
sharpness :: num II determines the shape of the sigmoid
sharpness $=2$
clip :: generator
clip v
$=$ low, if $v<=$ low
$=v \quad$, if low $<v<=$ high
$=$ high, if high $<v$

## Some well-know elements

```
perceptron_b :: element
perceptron_b weights = sgn . wsum weights || For binary values
perceptron_c :: element
perceptron_c weights = sigmoid . wsum weights || For continuous values
hopfield_b :: element
hopfield_b weights = hard_limit . wsum weights || For binary values
hopfield_c :: element
hopfield_c weights = clip . wsum weights II For continuous values
```

Functions to trace the evolution of the state and the learning of a net
evolution :: weights_net -> [input_state] $->$ [target_state] $->$ [input_state] $->$ [char]
evolution initial_weights samples targets inputs
$=$ lay [ lay [ show_state state | state <- compute_state weights input] | input <- inputs ] ++
show_w_net weights
where weights = leam Initial_weights samples targets

## II trace take 4 arguments:

\# - $f \quad$ : the function to be leamed
\# - size : the size of the network to be generate
I -m : the number of leaming cycles between status reports
$\|-\boldsymbol{n} \quad$ : the total number of leaming cycles to preform
trace $::$ (state $->$ state) $->$ [num] $->$ num $->$ num $\rightarrow$ [char]
trace f size m n
= "gain: " ++ show gain ++ " a: " + show a ++ "n" ++ concat [eval_perf (weights!i) (i*m) all_samples all_targets |i<-[0..reps]] ++
show_net (weights!reps)
where weights $=$ initial_net size: [learn (weights!i) (inputs!i) (targets!i) |i<-[0..]] inputs $\quad=$ chop $m$ input_seq targets $\quad=$ chop $m$ [f input | input <-input_seq] input_seq = all_samples ++ random_patterns dimension all_samples = patterns dimension all_targets $=$ [f input | input <- all_samples] reps $\quad=n$ div $m$ dimension = size! 0
eval_perf :: weights_net $\rightarrow$ num $->$ [state] $->$ [state] $->$ [char]
eval_pert weights $n$ samples targets
= "score after " ++ show n ++ " learning cycles: " ++ showfloat 0 (100 * score) ++ " \% ${ }^{\prime \prime} n^{n}++$ show_sot samples outputs targets
where score = sum [1|i<-[0..\#samples-1];outputs!i=targets $!$ i]/(\# samples)
outputs $=$ [compute_output weights input | input $<$ - samples]
show_sot $::$ [input_state] $->$ [output_state] -> [target_state] -> [char]
show_sot samples outputs targets
$=$ lay ["input:1t" ++ show_state (samples!i) ++
"toutput:1t" ++ show_state (outputsli) ++
"ttarget:lt" ++ show_state (targets!i) | i <- [0..\#samples-1]]
show_state $::$ state $->$ [char]
show_state state $=\left.{ }^{\prime \prime}\right|^{\prime \prime}+$ [show_elem elem | elem <- state] ++ "|"
show_w_net :: weights_net -> [char]
show_w_net net = layn [show_w_layer layer | layer <- net]
show_w_layer :: weights_layer -> [char]
show_w_layer layer = lay [show_w_vector vector | vector <- layer]
show_w_vector $::$ weights $\rightarrow$ [char]
show_w_vector vector $=$ concat ["It" ++ showfloat 2 elem| elem <- vector]

## Some list manipulation functions

update :: num -> * $\rightarrow$ [**] -> [*]
update 0 new (a:x) $=$ new:x
update $n$ new $(a: x)=a$ update $(n-1)$ new $x$

```
remove :: num -> ["] -> ["]
remove \(0(a: x)=x\)
remove \(n(a: x) \quad=a:\) remove \((n-1) x\)
to_limit::["]->["]
to_limit ( \(a: b: x\) )
\(=[a] \quad\), if \(a=b\)
= \(a\) : to_limit (b:x) , otherwise
scan_list :: num -> * -> ["] -> num
scan_list \(\mathrm{i} x\) (y:rest)
\(=\mathbf{i} \quad\), if \(x=y\)
\(=\) scan_list \((i+1) \times\) rest , otherwise
chop :: num -> [*] -> [[*]]
chop \(n\) seq \(=\) take \(n\) seq : chop \(n\) (drop \(n\) seq)
normalise :: vector \(->\) vector
normalise vector
\(=\) [element / I | element <- vector] , if I>0
= error "cannot normalise 0 -vector" , otherwise
    where \(I=\) length vector
length :: vector -> num
length vector \(=\) sqrt (sum [element^2 \(\mid\) element \(<-\) vector])
patterns :: num -> [state] \| Generates all binary patterns of length \(n\)
patterns n
= pats n [[ ]]
    where pats 0 states \(=\) states
    pats \(n\) states
    \(=\) pats (n-1) ([low : state | state <- states] ++ [high: state \| state <- states] )
random_patterns :: num -> [state] || Generates a list of random states
random patterns dimension
\(=\mathfrak{f}\)
    where \(\mathrm{fn}=\) pats!(n mod limit): f (random n ) || Pseudo random
        pats \(=\) patterns dimension
        limit \(=2^{\wedge}\) dimension
random_input :: num -> [state]
random_input dimension = chop dimension (random_values dimension)
random_values :: num -> [value]
random_values \(n\)
= value : random_values (random n)
    where value
    \(=\) high, if \(n \bmod 2=1\)
    = low , otherwise
```

Some functions to be learned

```
xor_fun :: input_state -> output_state
```

xor_fun inputs
$=[$ high $]$, if \# (filter (=high) inputs) $=1$
$=[$ low $]$, otherwise
and_fun :: input_state -> output_state
and_fun inputs
$=$ [high] , if \# (filter (=low) inputs) $=0$
$=[$ low $]$, otherwise
or_fun :: input_state -> output_state
or_fun inputs
$=$ [high] , if \# (filter (=high) inputs) $>=1$
$=$ [low] , otherwise
majority_fun :: input_state -> output_state
majority_fun inputs
$=[$ high $]$, if \# (filter (=high) inputs) $>$ \# inputs $/ 2$
$=[$ low $]$, otherwise

## C. 2 A feedback network to solve the $\mathbf{N}$-queens problem

```
high \(=1.0\)
low \(=0.0\)
queen \(=\) high
empty \(=\) low
nr_constant_neurons \(=\mathbf{0}\)
queen_generator \(v \quad=\operatorname{clip}(v+1 / 3)\)
queen_element weights \(\quad=\) queen_generator.wsum weights
compute_state = feed_back_asyn queen_element
compute_output weights input = select_output (compute_state weights input)
show_net = show_w_net
show_elem value
\(=\) ' Q ' , if \(v>0.9\)
\(=\) ' \(q\) ' , if \(v>0.7\)
\(={ }^{\prime *}\) ' , if \(v>0.5\)
\(=\) ' + ' , if \(v>0.3\)
\(=\) '.' , if \(v>0.1\)
\(={ }^{\prime}\), otherwise
    where \(v=\) (value - low) / (high - low)
initial_net \([n]=\left[\left[\left[\right.\right.\right.\) matrix_value \(\left.\left.\left.i j n \mid i<-\left[0 .\left(n^{\wedge} 2-1\right)\right]\right] \mid j \ll\left[0 . .\left(n^{\wedge} 2-1\right)\right]\right]\right]\)
```

```
matrix_value p1 p2 n
\(=1 \quad\), if \(\mathrm{p} 1=\mathrm{p} 2\)
\(=-1 / 2\), if \(i=k \vee j=1 \vee\) abs \((i-k)=a b s(j-1)\)
\(=0\), otherwise
    where ( \(\mathrm{i}, \mathrm{j}\) ) \(=\) coordinates p 1 n
                        (k, I\()=\) coordinates p 2 n
coordinates :: num -> num \(\rightarrow\) (num,num)
coordinates field_number board_size
= (row,col)
    where row = field_number div board_size
    col = field_number mod board_size
empty_board :: num -> input_state
empty_board \(n=\operatorname{rep}\left(\boldsymbol{n}^{\wedge} 2\right)\) empty
test :: num \(->\) [num] \(->\) [char]
test \(n\) positions
= monitor (compute_state (initial_net [n]) (place_queens positions (empty_board n)))
place_queens [] board = board
place_queens (p:rest) board = place_queens rest (update \(p\) queen board)
monitor :: [output_state] -> [char]
monitor outputs = concat [show_board board | board <- outputs]
show_board :: state -> [char]
show_board state
= border++lay [show_state row | row <- take n (chop n state)]++border
    where border = "+" ++ rep n '-' ++ "+ln"
    \(\mathrm{n} \quad=\) entier (sqrt (\# state))
```


## References

Aho, A.V., Sethi, R., Ullman, J.D. (1986). Compilers - Principles, Techniques and Tools. Addison-Wesley series in Computer Science.

Amit, D.J., Gutfreund, H., Sompolinsky, H. (1985). Storing infinite numbers of patterns in a spin-glass model of neural networks. Physical Review Letters, 55, pp 1530-1533.

Amit, D.J., Gutfreund, H., Sompolinsky, H. (1985). Spin-glass models of neural networks. Physical Review A. 32(2), pp 1007-1018.

Andres, V. (1990). The N queens problem: A neural Algorithm. Proccedings iASTED Int. Conf. on Modelling, Simulation and Optimization. May 2224, 1990. Montreal, Canada.

Backus, J.W. (1987). Can programming be liberated from the van Neumann style? A functional style and its algebra of programs. Communications of the $A C M, 21, \mathrm{pp}$ 613-641.

Bakel, S.J. van, (1990) A First Attempt in Polymorphic Type Assignment using Intersection Types. Techincal Report 90-9, May 1990, Department of Computer Science, University of Nijmegen, the Netherlands.

Barendregt, H.P. (1984) The lambda-calculus, Its Syntax and Semantixs (revised edition). Studies in logic and foundations of mathematics. 103, NorthHolland, Amsterdam.

Barendregt, H.P., Eekelen, M.C.J.D. van, Glauert, J.R.W., Kennaway, J.R., Plasmeijer, M.J., Sleep, M.R. (1987a). Term Graph Rewriting. Proceedings of Parallel Architectures and Languages Europe (PARLE), part II, Eindhoven, The Netherlands. Springer Lec. Notes Comp. Sci. 259. pp 141-158.

Barendregt, H.P., Eekelen, M.C.J.D. van, Glauert, J.R.W., Kennaway, J.R., Plasmeijer, M.J., Sleep, M.R. (1987b). Towards an Intermediate Language based on Graph Rewriting. Proceedings of Parallel Architectures and Languages Europe (PARLE), part II, Eindhoven, The Netherlands. Springer Lecture Notes on Computer Science. 259. pp 159-175.

Barendregt, H.P., Eekelen, M.C.J.D. van, Glauert, J.R.W., Kennaway, J.R., Plasmeijer, M.J., Sleep, M.R. (1988). Towards an Intermediate Language based on Graph Rewriting. Revised version. Journal of Parallel Computing 9 with selected papers of the conference on Parallel Architectures and Languages Europe (PARLE), Eindhoven, The Netherlands. North-Holland pp 163-177.

Beizer, B. (1983) Software Testing Techniques. Van Nostrand Reinhold electrical/computer science and engineering series. Van Nostrand Reinhold Company. New York. ISBN: 0-442-24592-0.

Bird, R. Wadler, P, (1988) Introduction to functional programming. Prentice Hall, New York.

Boute, R.T. (1988). System Semantics: Principles, Applications, and Implementation. ACM Transactions on Programming Languages and Systems, Vol. 10. No. 1. January 1988, pp 118-155.

Brus, T. Eekelen, M.C.J.D. van, Leer. M. van, Plasmeijer, M.J. (1987). Clean A Language for Functional Graph Rewriting. Proc. of the Third International Conference for Functional Programming Languages and Computer Architecturs (FPCA '87), Potland, Oregon, USA. Springer Lecture Notes on Computer Science 274, pp 364-384.

Burn, G.L., Peyton Jones, S.L., Robson, J.D. (1988). The spineless G-machine. Proceedings of the 1988 ACM Conference on LISP and Functional Programming.

Burstall, R.M., MacQueen, D.B., Sanella, D.T. (1980) Hope: An experimental Applicative Language. Proceedings of the 1980 LISP Conference. pp 136143.

Church, A. (1932 and 1933). A set of Postulates for the Foundation of logic. Annals. of Math. (2) 33 pp 346-366 and 34 pp 839-864.

Coolen, A.C.C. (1989). Inleiding Neurale Netwerken. Vakgroep Medische en Fysiologische Fysika, Rijksuniversiteit Uterecht.

Curry, H. B., Feys, R. (1958). Combinatory Logic. 1, North Holland.

DeMarco, T. (1987). Structured analysis and system specification. Yordon Press, New York.

Eekelen, M. C. J. D. van, Plasmeijer, M.J. (1986), Specification of rewriting strategies in Term Rewriting Systems. Proceedings of the workshop on Graph Reduction, Santa Fe, New Mexico. Springer Lec. Notes Comp. Sci. 279, 215-239.

Eekelen, M. C. J. D. van. (1988) Parallel Graph Rewrting, Some Contributions to its Theory, its Implementation and its Application. Ph. D. Thesis, University of Nijmegen.

Eekelen, M. C. J. D. van, Nöcker E. G. J. M. H., Plasmeijer, M. J. ,Smetsers, J. E. W. (1989) Concurrent Clean. Technical report 89-18, October 1989. University of Nijmegen. The Netherlands.

Ellacott, S.W. (1990) An analysis of the delta rule. Proceedings INNC (International Neural Network Conference), Paris, 1990. Volume II. Kluwer Academic Publisers. pp 956-990.

Fehr, E. (1989), Semantik von Programmiersprachen. Springer-Verlag. Berlin.
Field, A.J., Harrison, P.G. (1988), Functional programming. Addison-Wesley.
Gielen, C.C.A.M.(1988) Lecture notes caputcollege Neurophysics, Department of Medical and Biophysics, Nijmegen University. The Netherlands.

Gordon, M.J.C. (1979), The denotational description of programming languages. Springer-Verlag, New York.

Groningen. J.C. van. (1990). M. Sc. thesis. University of Nijmegen. To appear.
Hebb, D. O. (1949), The Organization of Behaviour. Wiley, New York.
Hecht-Nielsen, R. (1987a). Counter-propagation networks. IEEE First International Conference on Neural Networks, Vol II, pp 19-32.

Hecht-Nielsen, R. (1987b). Neurocomputer application. Proc. National Computer Conf. AFIPS. pp 239-244, 1987.

Hecht-Nielsen, R. (1988). Neurocomputing: picking the neural brain, IEEE Spectrum, Vol 25, no 3 pp 36-41, 1988.

Hoare, C.A.R. (1962). Quicksort. Computer Journal, 5, Nr 1, pp 10-15.
Hobbie, R.K. (1978). Intermediate Physics for Medicine and biology. Wiley New York.

Hopfield, J.J. (1982) Neural Networks and Physical systems with Emergent Collective Computional Abilities. Proc. Natl. Acad. Sci. USA, 79, pp 25542558.

Hopfield, J.J. and Tank, D.W. (1985). "Neural" Computation of Decisions in Optimization Problems. Biologocal Cybernetics, 52, pp 141-152.

Hudgkin, A.L. and Huxley A.F. (1952). "A quantitative description of membrane current and its application to conduction and excitation in nerve" Journal Physiology, 117: 500-544.

Hughes, J. (1989). Why Functional Programming matters. The computer journal, Vol 32, no 2, 1989 pp 98-107.

INNC 90. Proceedings INNC (International Neural Network Conference), Paris, 1990. Volume I \& II. Kluwer Academic Publisers. Dordrecht

Johnson, Th. (1984). Efficient compilation of lazy evaluation. Proceedings of the ACM SIGPLAN '84, Symposium on Compiler Construction. SIGPLAN Notices 19/6.

Johnson, Th. (1987). Compiling Lazy Functional Programming languages. Dissertation at Chalmers University, Götenborg, Sweden. ISBN 91-7032-280-5.

Joosten, S.M.M. (1989). The use of functional programming languages in software development. Dissertation Enschede. ISBN: 90-902729-7.

Kernighan, B.W., Ritchiw, D.M. (1978). The C programming language. Prentice-hall software series.

Kirkpatrick, S., Gelatt, C.D., Vecchi, M,P. (1983) Optimization by Simulated Annealing. Science, 220, pp 671-680.

Klop, J.W. (1987), Term Rewriting systems: a tutorial. Center for Mathematics and Computer Science, CWI Amsterdam, The Netherlands. Note CSN8701.

Kohonen, T. (1977), Associative Memory-A System Theoretic Approach. Springer, New York.

Kohonen, T. (1988a), Self-Organisation and Associative Memory. Springer, New York.

Kohonen, T. (1988b), An Introduction to Neural Computing, Neural Networks, 1, pp 3-16.

Kohonen, T. (1988c), Neural Networks, 1, pp 17-61.
Koopman, P.C. (1960). Neurotransmitters and their chemical derivates. A study on the relationship between structure and activity. Dissertation. R.C. University Nijmegen, The Netherlands.

Koopman, P.W.M. (1987). Interactive Programs in a Functional Language: A Functional Implementation of an Editor. Software-Practice and Experience, 17(9), pp 609-622.

Koopman, P.W.M., Nöcker, E.G.J.M.H.(1988) Compiling Functional Languages to Term Graph Rewrite Systems. Internal Report no. 88-1, Department of Computer Science, University of Nijmegen, The Netherlands.

Koopman, P.W.M., Eekelen, M.C.J.D. van, Nöcker, E.G.J.M.H., Smetsers, J.E.W., Plasmeijer, M.J. (1990). The ABC-machine: A Sequential Stackbased Abstract Machine For Graph Rewriting. Internal Report, University of Nijmegen, to appear in 1990.

Koopman. P.W.M., Rutten, L.M.W.J., Eekelen, van M.C.J.D. and Plasmeijer, M.J. (1990), Functional Description of Neural Networks, Proceedings INNC (International Neural Network Conference), Paris, 1990. Volume II. Kluwer Academic Publisers. pp 701-704.

Landin, P.J. (1964). The mechanical evaluation of expressions. Computer Journal. 6, pp. 308-320.

Ligner, G.T. (1975) A Mathematical Approach to Language Design. Proceedings of the second ACM symposium on priciples of programming languages, Palo alto.

Lippman, R.P. (1987). An Introduction to Computing with Neural Nets. IEEE ASSP Magazine April 1987. pp 4-22

Milner, R.A. (1978). Theory of Type Polymorphism in Programming. Journal of Computer System Sciences, 17, no 3. pp 348-375.

Minsky, M.L., Papert, S.A. (1969), Perceptrons. MIT, Cambridge (Mass).
Mycroft, A. (1984). Polymorphic type schemes and recursive definitions. Proc. of the 6th Int Conf on Programming. Springer Lecture Notes on Computer Science 167. pp 217-228.

Nöcker, E.G.J.M.H. (1989) The PABC Simulator v0.5. Implemenation Manual. Technical report no. 89-19, October 1989. Department of Informatics. University of Nijmegen. The Netherlands.

Nöcker, E.G.J.M.H. (1990) Strictness analysis and pattern matching. Technical report. University of Nijmegen. The Netherlands. To appear.

Peyton Jones, S.L. (1987). The implementation of functional programming languages. Prentice-Hall. ISBN 0-13-453325-9.

Plasmeijer, M.J., Eekelen, van M.C.J.D.(1991). Functional Programming and Parallel Graph Rewritting. To appear by Eddison Wesly. Currently available as preprint at the department of computer science, university of Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands.

Reif, F. (1965). Fundamentals of statistical and thermal physics.McGraw-Hill
Reichl, L.E. (1980), A Modern Course in Statistical Physics. E.A. London.
Ridgway, W.C. (1962). An adaptive Logic System with Generalizing Properties, PhD thesis, Standford Electronnics Labs, Stanford University, April 1962.

Rosenblatt,.F. (1958). The perceptron: A probalistic model for information storage and organization in the brain. Psychoanlytical Review, 65, pp 386408.

Rosenblatt, F. (1962), Principles of Perceptrons. Spartan, Washington.
Rumelhart, D.E., Hinton, G.E., Williams, R.J. (1986). Learning internal representations by error propagation. In Rumelhart \& McClelland: Parallel Distributed Processing: Explorations in the microstructure of Cognition. Vol 1: Foundations. MIT Press.

Schadé. J.P. (1984). Onze hersenen, Aula pocket 721, Het Spectrum, Utrecht/Antwerpen.

Schönfinkel, M. (1924). Über die Bausteine der mathematischen Logik. Math. Annalen 92. pp 305-316.

Simpson, P.K. (1990) Artificial Neural Systems. Pergamon Press. ISBN 0-08-037894-3

Smetsers, J.E.W. (1987). Parallel graph reduction on a distributed machine. M. Sc. thesis, Philips Research Laboratories, University of Nijmegen, October 1987.

Smetsers, J.E.W. (1989). Compiling Clean to Abstract ABC-Machine Code. Technical Report 89-20. Department of Computer Science, University of Nijmegen, the Netherlands.

Smetsers, J.E.W. Nöcker, E.G.J.M.H., Koopman, P.W.M., Plasmeijer, M.J., Eekelen, van M.C.J.D.(1990). The parallel ABC-machine. To appear.

Stoy J.E. (1977). Denotational Semantics: The Scott-Strachy Approach to Programming Language Theory. Cambridge (Massachusetts), MIT Press.

Tanenbaum, A.S. (1984). Structured computer organization, second edition. Prentice-Hall.

Tank, D.W., Hopfield, J.J. Collective Computation in Neuronlike Circuits. Scientific American.

Turner, D. A., (1976). SASL Language Manual; (1979) SASL Language Manual, "combinators" version. University of St Andrews.

Tumer, D. A., (1979). A New Implementation Technique for Applicative Languages. Software-Practice \& Experience, 9(1), pp 31-49.

Turner, D. A., (1985). Miranda: A non-strict functional language with polymorphic types. Proc. of the conference on Functional Programming Languages and Computer Architecture, Springer Lecture Notes on Computer Science 201.pp 1-16.

Weijers, G.A. (1990), ABC-machine implementation Guide, University of Nijmegen. To appear.

Widrow, B. (1962). Generalization and information storage in networks of Adaline 'neurons', in self-organzing systems 1962, Yovits, M.C., Jacobi, G.T. and Goldstien, G. eds, Spartan Books, Washington, DC, pp 435-461.

Widrow, B. and Streams, S.D.(1985), Adaptive signal processing, Prentice hall, Englewood Cliffs, N,J.

Widrow, B.,Winter, R. (1988) Neural nets for adaptive filtering and adaptive pattem recognition, IEEE Computer, March 1988.

Wirth, N. (1971) The programming language PASCAL. Acta Informatica 1, 1971. pp 35-63.

## Index

## Keywords and function names

_cycle 100
_driver 100
_mf 100
$\alpha$-conversion 60
A-stack 41, 46, 151, 154
ABC-machine 16, 35, 80, 82, 151, 153
ABC-assembler 50, 151, 168, 175
ABC-code 79, 82, 175
abstract program translations 57, 79, 168
abstract syntax tree 57
abstraction 62,70, 74
adaline $117,125,128$
adapt_weights 120
Adaptive Linear Combiner 125
add_args 159, 161, 168
addI $160,166,169$
additional rule alternative 83, 99
ALC 125
alt_entry 93, 177, 178
alt_label 184
and_fun 191
annealing 186
AP 82, 84, 101
apply entry $84,92,101,103$
apply_label 184
argument 98
asp 93
assemble 169
AST 57, 73, 89, 168, 175
axon 112
B 63, 69, 74
B-stack 42, 47, 103, 151, 154
$\beta$-reduction 60
backpropagation 123, 130, 137
bind 183
bind_args 183
bindings 93
bool_symbolid 184
boot 49, 167
bound 59
Br 168
Br false 168
Br_true 168
bracket abstraction 60
bring_in_mf 183
build_arg 98, 181
build_args 98, 181
C 63, 69, 74
C-stack 43, 152, 155
calling conventions 83,103
chop 190
Church's $\delta$ rule 31
Clean 16, 79, 80, 85, 175
clean_up 184
clip 116, 188
combinator 62
combinatorial logic 60
compilation scheme 90
compile 90, 175
computational model 30
cons_descriptor 91, 176
constant neurons 118
constant_neurons 118, 186
contractum 81, 85, 96, 98, 179
copying 86
counter propagation 136
create $44,159,161,168$
Curried 82, 101
cycle 96
cycle_lab 184
cycle_nodes 96, 180
cyclic dependencies 117,118
decI $160,166,169$
del_args 159, 161, 168
delta rule 120, 128, 131, 134
desc-id 39
descriptor $40,91,153,157,168,176$
descriptor store $40,153,157$
difference vector 136
disturb 186
dump 159, 161, 168
element 117
empty_bindings 183
energy 139
eq_desc $159,161,168$
eq_desc_arity 46, 159, 161, 168
eq_symbol 159, 162
eqB $159,162,168$
eqB_a 159, 162, 168
eqB_b $159,162,168$
eqI $159,162,168$
eqI_a $46,159,162,168$
eqI_b 159, 162, 168
error 121
escape_false 95, 183
eval_layer 118, 186
eval_neurons 187
eval_one_by_one 187
eval_perf 189
evolution 188
extensional equal 64
extensional equality 30
feed_back_asyn 119, 186
feed_back_syn 119, 186
feed_forward 118, 186
feedback networks 118, 138, 142
fetch_cycle 49, 167
FGRS 80
fill 44, 97, 159, 162, 168
fill_a 159, 162, 168
fill_code 92, 177
fill_node 97, 180
fill_nodes 96, 180
fill_strict 97, 181
fill_with_mf 100, 177
fillB $159,163,168$
fillB_b 160, 163, 168
fill $160,163,168$
filli_b $160,163,168$
fixed point 66
fixed point combinator 66
flow control 47
free 59
fun_descriptor 91, 176
Functional Graph Rewriting System 80
functional strategy 80,85
G-machine 51
garbage 81
gen_type_error 100,177
generalized delta rule 120, 127, 129, 131
generator 115
get_args 183
get_desc_arity $160,163,168$
get_lhs 182
get_node_arity $160,163,168$
get_rhs 182
gradient descent 121
Graph manipulation 44
Graph rewriting systems 16,80
graph store $39,152,157$
GRS 80
gtI 160, 166, 169
halt 48, 160, 163, 168
Hamming network 136
hard_limit 116, 188
head normal form 71
Hebb 120, 138, 140
hidden elements 121
hierarchical machine description 19
higher order 82
hnf 71
Hodgkin Huxley equations 112
hopfield 117
hopfield_b 188
hopfield_c 188
I 62, 69, 74
I/O-channels 153, 158
in situ 12
incl $160,167,169$
indirection node 86
initial_net 117
initial_status 184
inner_product 187
input-output channel 43
instr-id 37, 39, 43
instruction 22
instruction level 20, 22
instruction set 44, 159
instruction_cycle 24
int_symbolid 184
is_strict 182
Ising spin model 142
jmp 47, 160, 163, 168
jmp_eval $160,163,168$
jmp_false 47, 160, 164, 168
jmp_true 47, 160, 164, 168
jsr 23, 48, 160, 164, 168
jsr_eval 48, 160, 164, 168
K 62, 69, 74
$\lambda$-calculus 16, 59
Label 168
last entry 99
law of abstraction 62
lazy evaluation 80
learn 120, 187
Learning 115, 117, 120, 122
Learning Vector Quantization 128, 134
left hand side 81
Length 81, 88, 104, 190
lhs 81
linear separable 132
list 65
listing 173
LMS 124
Local definitions 65
local function 75
locus of control 43
ltI 160, 167, 169
LVQ 134, 146
Mac-1 25
Mac-1 instructions 29
machine state 20
madaline 126
majority_fun 191
match 94,178
match_arg 95, 178
match_args 94, 178
matching 84,93
McCulloch-Pitts 114, 120, 138, 145
median 10
membrane permeability 112
micro-instruction set 38
micro-instructions 20, 21
Miranda 11, 16
Miranda system 30
Miranda interpreter 30
muli $160,167,169$
N -flop 134, 140
$N$-queens 141, 143, 185, 191
neuron 112
neuron activity 113
neurotransmitters 112
nf 81
nfib-number 107
nil 65
no_op $160,164,168$
node $39,152,156$
node entry 84, 92, 103
node-id 39
node_label 184
noise 186
normal form 81
normal order 61
normalise 190
Operational Machine Description 19, 35
optimization rule 70
optimization rules 64,75
optimizations 103
or_fun 191
overwriting 85,86
P65, 74
partially parametrized functions 30
pattern 81
patterns 190
perceptron 117, 126, 127, 145
perceptron_b 188
perceptron_c 188
pop_a $46,160,164,168$
pop_args 183
pop_b 160, 164, 168
potential 115
prepare_args 92, 177
print 48, 160, 164, 169
print_symbol 48, 160, 164, 169
program 42
program counter $24,43,153,158$
Program execution 24
Program Store 153, 158
prototype 11
push_a $46,160,164,169$
push_ap_entry 160, 165, 169
push_arg 160, 165, 169
push_arg_b 160, 165, 169
push_args $45,160,165,169$
push_args_b 160, 165, 169
push_b $160,165,169$
pushargs 183
pushB 160, 165, 169
pushB_a $160,165,169$
pushi $160,165,169$
pushI_a $45,160,166,169$
quicksort 10
random 187
random_input 190
random_patterns 190
random_values 190
record 184
record_args 183
Recursion 66
redex $59,71,81$
redirection $81,85,99,182$
reduce_arg 93, 177
reduce_strict_args 92, 177
reduction 59,71
reduction_label 184
remove 190
repl_args $160,166,169$
repl_args_b $160,166,169$
Rewrite alternative 93
rewriting 85
rhs 81
right hand side 81
mf 81
mf_lab 184
root 99, 182
root normal form 81, 98
rtn 23, 48, 160, 166, 169
rule alternative 83
rule alternative entry 84
run-time system 83, 100
S 62, 69, 74
SASL 61
scan_list 190
select_output 118, 186
set_entry $160,166,169$
sgn 116, 188
shared_nodes 96, 179
sharpness 188
show_sot 189
show_state 189
show_w_layer 189
show_w_net 189
show_w_vector 189
sigmoid 116, 129, 188
sigmoid' 188
Sorting 9
spike 112
stack-frame 83
state $94,114,175$
status 94,175
strict $81,89,92,97,103$
subI $160,167,169$
substitution 59
symbol_id 184
synapse 112, 113, 115, 117
syntactical equal 30
tail recursion 99
Term Rewriting System 80
threshold 116, 123
to_limit 190
trace 189
translate 170
TRS 80
type_error 100, 184
U 65, 69, 74
unravelling 86
update 189
update_a $46,160,166,169$
update_astack 183
update_b $160,166,169$
value 114
vector 134
weights 115
wsum 115, 121, 127, 187
xor_fun 191
Y $66,69,74$

## Samenvatting Functionele programma's als uitvoerbare specificaties

Dit proefschrift behandelt het gebruik van functionele programmeertalen voor het opstellen van specificaties die door een computer kunnen worden uitgevoerd.

Een specificatie is een zeer nauwkeurige beschrijving in een geschikte beschrijvingstaal. Een bruikbare specificatie is duidelijk, foutloos en voldoende compleet. De gebruikte beschrijvingstaal mag geen dubbelzinnigheden bevatten en moet krachtig genoeg zijn om compacte specificaties mogelijk te maken. In de informatica worden specificatics gebruikt om eigenschappen en gedragingen van machines en programma's vast te leggen. Een specificatie in een functionele programmeertaal is niet alleen een beschrijving van de eigenschappen, maar geeft ook een algoritme: een voorschrift om die eigenschappen te verwezenlijken.

In traditionele programmeertalen bestaan programma's uit een lange reeks bevelen. Deze talen heten daarom imperatieve talen. Het gevolg van een bevel is afhankelijk van de eerder uitgevoerde bevelen. Dergelijke programma's zijn daardoor vaak moeilijk te begrijpen en het is niet eenvoudig hun correctheid aan te tonen.

In een functionele taal bestaat een programma uit een reeks functicdefinities en een beschrijving van de gewenste oplossing in termen van deze functies. De functies geven eigenschappen van het beschreven object. Zij worden gebruikt om de beschrijving van het resultaat in de meest eenvoudige vorm te krijgen. De beschrijving moet zo gekozen worden dat niet alleen de goede oplossing beschreven wordt, maar ook een manier om die oplossing te berciken. Doordat tijdens het vereenvoudigen de betekenis van de beschrijving behouden blijft zijn programma's in een functioncle taal duidelijk en is de correctheid relatief eenvoudig te bewijzen. Dergelijke programma's zijn bovendien vaak erg kort en bondig.

Om vertrouwen te krijgen in de correctheid van de beschrijving en de geschiktheid van het beschreven produkt wordt vaak een prototype gemaakt. Een prototype is een gedeeltelijke, relaticf snel te maken implementatie van het beschreven produkt. Het prototype geeft inzicht in het gedrag van het beschre-
ven produkt, maar heeft een aantal beperkingen ten opzichte van het definitieve produkt. Bij een prototype van een programma gaat het primair om een bepaald gedrag en bijvoorbeeld niet om de snelheid of de manier waarop het programma wordt aangeroepen.

Het is aantrekkelijk programmeertalen als beschrijvingstaal te gebruiken. Programmeertalen zijn doorgaans zo goed gedefiniëerd dat het duidelijk is wat de beschrijving precies betekent. De implementatie van de programmeertaal kan gebruikt worden om een aantal aspecten van de beschrijving automatisch te controleren. Een compiler controleert bijvoorbeeld of er geen taalfouten gemaakt zijn en of alle gebruikte termen gedefiniëerd zijn. Bovendien is de specificatie een programma dat door een computer kan worden uitgevoerd, de beschrijving is dus zijn eigen prototype.

Ondanks deze voordelen worden programmeertalen zelden als beschrijvingstaal gebruikt: zij zijn meestal niet krachtig genoeg om korte en duidelijke beschrijvingen mogelijk te maken. Wanneer we toch zo'n programmeertaal zouden gebruiken dan wordt de beschrijving zó lang dat de duidelijkheid verloren gaat. Bovendien verdwijnt op die manier het verschil tussen beschrijving en produkt.

Functionele programmeertalen zijn hoog-niveau-programmeertalen met een grote uitdrukkingskracht, ze lijken daarom geschikter als beschrijvingstaal dan de traditionele imperatieve talen. In dit proefschrift worden enkele beschrijvingsmethoden geïntroduceerd om de geschiktheid van functioncle programmeertalen als beschrijvingstaal te onderzoeken. Aan de hand van enige voorbeelden worden deze beschrijvingen vergeleken met traditionele beschrijvingen. De bruikbaarheid van deze beschrijvingsmethoden voor uitgebreide specificaties wordt onderzocht door ze toe te passen op enkele voorbeelden uit het dagelijks leven van een informaticus.

In hoofdstuk 2 wordt een machine-beschrijvingsmethode geıntroduceerd die bestaat uit meerdere lagen. In de onderste laag worden de onderdelen van de machine, bijvoorbeeld de geheugens, beschreven. De volgende laag beschrijft de instructies van de machine door het effect van iedere instructie op de onderdelen van de machine aan te geven. De verkregen beschrijving wordt vergeleken met een traditionele beschrijving. De machine-beschrijving volgens de nieuwe methode blijkt iets langer te zijn dan de oude specificatie. Dit wordt veroorzaakt door de overzichtelijke manier van noteren die we onszelf hebben opgelegd. De functionele beschrijving blijkt erg duidelijk te zijn en is bovendien iets directer dan de traditionele beschrijving. Het is tevens een bruikbaar prototype.

Graafherschrijfsystemen zijn een zeer geschikt berekeningsmodel voor de implementatie van functionele programmeertalen. Een programma in een functionele programmeertaal worden eerst vertaald naar een equivalent programma in de graafherschrijftaal Clean. Een Clean-programma wordt vervolgens vertaald naar een programma voor een abstracte, imperaticve graafherschrijfmachine: de ABC-machine. Deze abstracte machine is ontworpen om aan te kunnen geven hoe een Clean-programma wordt uitgevoerd op imperatieve machines zonder last te hebben van de beperkingen van een concrete machine.

Als laatste stap wordt een ABC-programma vertaald naar code voor een concrete machine. Het blijkt dat via deze tussenstappen uitstekende implementaties van functionele programmeertalen verkregen worden. De beschrijvingsmethode uit hoofdstuk 2 wordt in hoofdstuk 3 gebruikt voor de beschrijving van de ABC-machine. De gegeven specificatie is de enige machine-beschrijving die beschikbaar is en voldoet goed in het gebruik.

In hoofdstuk 4 wordt de geschiktheid van functionele talen voor het beschrijven van vertalingen van programmeertalen onderzocht. De te transformeren taal wordt gerepresenteerd door een datastructur in de functionele beschrijvingstaal. De vertaling is een functie, die de gewenste transformatie op de datastructuren uitvoert. Opnieuw voldoet de specificatie in een functionele programmeertaal goed in vergelijking met andere beschrijvingen.

De vertaling van de graafherschrijftaal Clean naar instructies voor de in hoofdstuk 3 beschreven ABC-machine wordt in hoofdstuk 5 behandeld. Deze beschrijving dient als ontwerp voor de echte vertaler. Het uiteindelijke produkt bevat veel extra optimalisaties. Het is mogelijk om al die optimalisaties in de specificatie op te nemen, maar dan zou het basisschema van de vertaling onvoldoende duidelijk naar voren komen. De beschrijvingen uit hoofdstuk 3 en 5 zijn samen gebruikt om Clean programma's uit te voeren en om hun gedrag te bestuderen voordat de beschreven producten gemaakt werden.

In hoofdstuk 6 worden een aantal kunstmatige neurale netwerken beschreven in een functionele taal. Deze beschrijvingen zijn een zinvolle aanvulling op de traditionele wiskundige beschrijvingen. Wiskundige specificaties zijn, vaak niet complete, beschrijvingen van eigenschappen van deze neurale netwerken, maar geven niet aan hoe dit gedrag op een efficiënte manier bereikt kan worden. De ontwikkelde beschrijvingen laten de overeenkomsten en verschillen tussen de netwerken duidelijk zien en zijn gebruikt als prototype om een aantal niet wiskundig te bepalen eigenschappen te onderzoeken.

Als conclusie kunnen we stellen dat een functionele programmeertaal met succes is gebruikt voor een aantal uitgebreide beschrijvingen die de vergelijking met andere specificaties glansrijk kunnen doorstaan.

Voor alle in dit proefschrift gegeven beschrijvingen is de functionele programmeertaal Miranda gebruikt. Deze taal is niet speciaal als beschrijvingstaal ontworpen, maar is desalniettemin goed bruikbaar als beschrijvingstaal. Op twee punten deden zich problemen voor. Ten eerste was het, door het ontbreken van herschrijf-semantiek, niet mogelijk om de instructies in het geheugen van de in hoofdstuk 2 en 3 beschreven machines te laten zien. Ten tweede maakten het gebruikte typeringssysteem het noodzakelijk om enkele, overigens overbodige, constructoren te introduceren in de datastructuren die voor de beschrijving van vertalingen in hoofdstuk 4 en 5 werden gebruikt.

## Curriculum vitae

3 juli 1957 geboren te Nijmegen.
1970-1976 Atheneum-B aan het St. Odulphus Lyceum te Tilburg.
1976-1980 Kandidaats examen Natuurkunde (N4: hoofdvakken Natuurkunde en Scheikunde met bijvak Wiskunde) aan de Katholieke Universiteit Nijmegen (KUN).

1980-1983 Doctoraal examen Experimentele Natuurkunde met bijvakken capita uit de Natuurkunde en Informatica aan de KUN. Het afstudeerwerk stond onder begeleiding van Dr. P.I.M. Johannesma en betrof tweede-orde-signaalverwerking. Dit werk vond plaats op de afdeling Medische- en Biofysica onder leiding van Prof. Dr. J. Eggermont.

1984-1990
Wetenschapelijk medewerker/universitair docent/toegevoegd onderzoeker bij de sectie/discipline/vakgroep Informatica van de Faculteit Wiskunde en Natuurwetenschappen/Faculteit Wiskunde en Informatica aan de Katholieke Universiteit Nijmegen. Aanvankelijk was ik werkzaam bij de afdeling Machinearchitectuur en Bedrijfssystemen onder verantwoordelijkheid van Prof. Dr. Ir. R.T. Boute. Vervolgens heb ik gewerkt in de vakgroep Theoretische Informatica en Berekeningsmodellen bij Prof. Dr. H.P. Barendregt. Dit proefschrift is voltooid terwijl ik werkzaam was bij de afdeling Parallele Systemen en Berekeningsmodellen onder leiding van Dr. Ir. M.J. Plasmeijer.



[^0]:    ${ }^{1}$ For instance, it can be shown that the informally specified assignment rule for ALGOL 60 can fail to hold in six ways [Ligler 75].
    ${ }^{2}$ The clumsy handling of the median is necessary in order to obtain termination for a sequence with some equal elements.

[^1]:    ${ }^{3}$ Miranda is a trademark of research software Ltd.

[^2]:    ${ }^{1}$ The Miranda system version 2.009 (13 November 1989) of Research Software Ltd. is referenced in this thesis as 'the Miranda interpreter'. All timings are done on a SUN-3/280.

[^3]:    ${ }^{1}$ Note: not all pointers are drawn; The symbol names and labels in the nodes are actually pointers in the descriptor store and the program store. Using arrows for these pointers would make the picture unreadable.

[^4]:    ${ }^{1}$ It is of course possible and more efficient to incorporate these improvement rules in the abstraction algorithm. For the sake of clarity, abstraction and optimization of the expression by the improvement rules are treated as separate algorithms here. Moreover, the combined algorithm is hard to specify elegantly in this formalism.

[^5]:    ${ }^{2}$ This name is due to Turner, 'unpair' or 'unpack' seems to be a better name.

[^6]:    ${ }^{1}$ In this chapter ZF-expressions are used to specify matrix and vector like operations, instead of the equivalent map's, zip's and transposes. The given definition resembles the mathematical expression

    $$
    \sum_{i=0}^{N-1} \text { weigths }_{i} \times \text { inputs }{ }_{i}
    $$

    more than the equivalent

    > sum (map2 (*) weights inputs)

[^7]:    $\qquad$

