
CAL Language Report

Specification of the CAL actor language



Johan Eker
Jörn W. Janneck

language version 1.0 — document edition 1

ERL Technical Memo UCB/ERL M03/48
University of California at Berkeley
December 1, 2003

Contents

Contents	2
1 Introduction	6
1.1 Actors and actor composition	6
1.2 Language design: goals and principles	8
1.3 Platform independence and compatibility	9
I Language description	12
2 Introductory remarks	13
2.1 Lexical tokens	13
2.2 Typographic conventions	15
2.3 Syntax rules	15
2.4 Notational idioms	15
3 Structure of actor descriptions	17
3.1 Namespaces and imports	18
3.2 Time	19
4 Data types	21
4.1 Objects, variables, and types	21
4.2 Type formats	23
4.3 Required types	24
4.4 Structured objects and their types	25
4.5 Mutable objects and their types	26
4.6 Type framework	27
5 Variables	28
5.1 Variable declarations	28
5.1.1 Explicit variable declarations	29
5.2 Variable scoping	30

6	Expressions	35
6.1	Literals	36
6.2	Variable references	36
6.2.1	Old variable references	36
6.3	Function application	37
6.4	Field selection	38
6.5	Indexing	39
6.6	Operators	39
6.7	Conditional expressions	39
6.8	Introducing a local scope	40
6.9	Closures	40
6.9.1	Lambda-expressions and function closures	40
6.9.2	Proc-expressions and procedure closures	41
6.9.3	Function and procedure declarations	42
6.10	Comprehensions	42
6.10.1	Simple collection expressions	43
6.10.2	Comprehensions with generators	43
6.11	Type assertion	46
7	Statements	48
7.1	Assignment	48
7.1.1	Simple assignment	49
7.1.2	Field assignment	49
7.1.3	Assignment with indices	49
7.1.4	Assigning to and from mutable variables	49
7.2	Procedure call	51
7.3	Statement blocks (begin ... end)	52
7.4	If-Statement	52
7.5	While-Statement	52
7.6	Foreach-Statement	53
7.7	Choose-Statement	54
8	Actions	57
8.1	Input patterns, and variable declarations	58
8.1.1	Single-port input patterns	59
8.1.2	Multiport input patterns	60
8.1.3	Scoping of action variables	61
8.2	Output expressions	62
8.3	Delays	65
8.4	On action selection: guards and other activation conditions	65
8.5	Initialization actions	66

9	Action-level control structures	68
9.1	Action tags	69
9.2	Action schedules	70
9.2.1	Finite state machine schedules	70
9.2.2	Regular expression schedules	71
9.3	Priorities	73
II	Semantics	75
10	Actor model	76
10.1	Preliminaries	76
10.2	Time systems	77
10.3	Actor transition systems	79
III	Appendices	81
A	CAL language syntax	82
A.1	Actor	82
A.2	Expressions	83
A.3	Statements	84
A.4	Actions	85
A.5	Action control	86
B	Keywords	87
C	Basic runtime infrastructure	90
C.1	Operator symbols	90
C.2	Basic data types and their operations	91
C.2.1	Collection[T]—collections	91
C.2.2	Seq[T]—sequences	93
C.2.3	Set[T] < Collection[T]—sets	94
C.2.4	List[T] < Collection[T], Seq[T]—lists	95
C.2.5	Map[K, V]—maps	96
C.2.6	Number—numbers	97
C.2.7	Character—characters	97
C.2.8	String < List[Character]—strings	97
C.2.9	Boolean—truth values	97
C.2.10	Null—the null value	98
	Index	100
	Bibliography	111

First, we want to establish the idea that a computer language is not just a way of getting a computer to perform operations but rather that it is a novel formal medium for expressing ideas about methodology. Thus, programs must be written for people to read, and only incidentally for machines to execute. Second, we believe that the essential material to be addressed by a subject at this level is not the syntax of particular programming-language constructs, nor clever algorithms for computing particular functions efficiently, nor even the mathematical analysis of algorithms and the foundations of computing, but rather the techniques used to control the intellectual complexity of large software systems.

[...]

Underlying our approach to this subject is our conviction that “computer science” is not a science and that its significance has little to do with computers. The computer revolution is a revolution in the way we think and in the way we express what we think. The essence of this change is what might best be called procedural epistemology—the study of the structure of knowledge from an imperative point of view, as opposed to the more declarative point of view taken by classical mathematical subjects. Mathematics provides a framework for dealing precisely with notions of “what is.” Computation provides a framework for dealing precisely with notions of “how to”

Harold Abelson, Gerald Jay Sussman
Structure and Interpretation of Computer Programs [2]

What a thing means is simply what habits it involves [...] there is no distinction of meaning so fine as to consist in anything but a possible difference in practice.

Charles S. Peirce [13]

Chapter 1

Introduction

This report describes CAL, an actor language created as a part of the Ptolemy II project [1] at the UC Berkeley. It is intended primarily as a repository for technical information on the language and its implementation and contains very little introductory material. After a short motivation, we will outline the goals and the guiding principles of the language design. We will also give a short outline of the actor model, and the context that the actors written in CAL are embedded into, describing the kinds of assumptions an actor may and may not, in general, make about it.

1.1 Actors and actor composition

Actors. The concept of actor as an entity that is composed with other actors to form a concurrent system has a rich and varied history—some important milestones include [6], [9], [3], [4], [5]. A formal description of the notion of actor underlying this work can be found in chapter 10, which is based on the work in [10] and [7]. Intuitively, an actor is a description of a computation on sequences of tokens (atomic pieces of data) that produces other sequences of tokens as a result. It has input ports for receiving its input tokens, and it produces its output tokens on its output ports.

The computation performed by an actor proceeds as a sequence of atomic steps called firings. Each firing happens in some actor state, consumes a (possibly empty) prefix of each input token sequence, yields a new actor state, and produces a finite token sequence on each output port.¹

Several actors are usually composed into a network, a graph-like structure (often referred to as a model) in which output ports of actors are connected to input ports of the same or other actors, indicating that tokens produced at those output ports are to be sent to the corresponding input ports. Such actor

firing

composition
of actors

¹The notion of actor and firing is based on the one presented in [10], extended by a notion of state in [7].

networks are of course essential to the construction of complex systems, but we will not discuss this subject here, except for the following observations:

- A connection between an output port and an input port can mean different things. It usually indicates that tokens produced by the former are sent to the latter, but there are a variety of ways in which this can happen: token sent to an input port may be queued in FIFO fashion, or new tokens may 'overwrite' older ones, or any other conceivable policy. It is important to stress that actors themselves are oblivious to these policies: from an actor's point of view, its input ports serve as abstractions of (prefixes of) input sequences of tokens, while its output ports are the destinations of output sequences.
- Furthermore, the connection structure between the ports of actors does not explicitly specify the order in which actors are fired. This order (which may be partial, i.e. actors may fire simultaneously), whether it is constructed at runtime or whether it can be computed from the actor network, and if and how it relates to the exchange of tokens among the actors—all these issues are part of the interpretation of the actor network.

decoupling
from commu-
nication
model

decoupling
from actor
scheduling

The interpretation of a network of actors determines its semantics—it determines the result of the execution, as well as how this result is computed, by regulating the flow of data as well as the flow of control among the actors in the network. There are many possible ways to interpret a network of actors, and we call any specific interpretation a model of computation—the Ptolemy project focuses on exploring the issues of models of computation and their composition, cf. [11, 12]. Actor composition inside the actor model that CAL is based on has been studied in [8].

communication
+ scheduling
= model of
computation

As far as the design of a language for writing actors is concerned, the above definition of an actor and its use in the context of a network of actors suggest that the language should allow to make some key aspects of an actor definition explicit. These are, among others:

- The port signature of an actor (its input ports and output ports, as well as the kind of tokens the actor expects to receive from or be able to send to them).
- The code executed during a firing, including possibly alternatives whose choice depends on the presence of tokens (and possibly their values) and/or the current state of the actor.
- The production and consumption of tokens during a firing, which again may be different for the alternative kinds of firings.
- The modification of state depending on the previous state and any input tokens during a firing.

Actor-like systems. It is often useful to abstract a system as a structure of cooperating actors. Many such systems are dataflow-oriented, i.e. they consist of components that communicate by sending each other packets of information, and whose ability to perform computation depends on the availability of sufficient input data. Typical signal processing systems, and also many control systems fall into this category.

Writing actors is hard. Writing an actor in a general-purpose programming language is of course possible, but most or all of the information that may be used to reason about its behavior is implicit in the program and can only be extracted using sophisticated analysis, if this is at all feasible.

Furthermore, actors often need to be run on different platforms. For instance, if actors are used in the design of an embedded system, they need to run in a modeling and simulation environment (such as Matlab or Ptolemy) as well as in the final product. Being able to use the same description of the actor functionality in both cases improves productivity and reduces the probability of errors.

1.2 Language design: goals and principles

Designing a programming language is an exercise in balancing a number of sometimes contradicting goals and requirements. The following were the ones that guided the design of CAL.

Ease of use. CAL is intended to be a programming language, not an intermediate format or a representation for automatically generated code. Since we want people to actually write code in it, the notation must be reasonably convenient to write, with consistent syntax rules, keywords, and structures. Because people make mistakes, it needs to be sufficiently redundant to allow effective error detection and localization, but simple and concise enough for frequent use, especially in frequently used areas.

Minimal semantic core. In spite of being a full-fledged programming language, we wanted to build CAL on a very small set of semantic concepts, for a number of reasons. First of all, being able to describe a large part of the full language through reductions to a smaller language makes the definition of language semantics much easier. From a practical perspective, this simplifies compiler construction—if there is a generic procedure that transforms any program into an equivalent program in the core language, then all it takes in order to compile the full language to any given platform is a code generator for the core language. This led to the design of a number of core languages of CAL (reflecting several levels of reduction) that much of the implementation of the language is based on, and which will be described in another place.

Focus and specificity. CAL is a domain-specific language that is aimed at providing a medium for defining actors. It was very important to draw a clear line between those pieces of functionality that were deemed to be part of an actor definition and those that were not. For example, in addition to clearly actor-specific structures such as actions and input/output patterns/expressions, expressions and statements were considered to be essential to defining an actor. On the other hand, there are many things that CAL explicitly does not contain, such as facilities for defining new types, concepts for connecting actors, or mechanisms to aggregate actors into composites. The fact that CAL is explicitly agnostic about these issues makes it possible to use the language in a wide variety of contexts, which may provide very different designs in those areas.

Implementation independence and retargetability. Even though our first target for CAL actors is the Ptolemy II platform, we want the language to be retargetable, in the following two senses: First, we would like to be able to take an actor written, say, for Ptolemy II and be able to compile it to some other platform, say to some C code that runs against a different API. Secondly, we would like to enable other people to embed CAL into entirely different, but still actor-like, contexts, which have different kinds of objects (and types), different libraries, different primitive data objects and operators. Here, we would not necessarily try to reuse the actor libraries written for other platforms (although some interesting subset might still be sufficiently generic to be reusable)—instead, we would reuse the CAL framework, i.e. its infrastructure such as parsers, transformers and annotators, verification and type checking, code generation infrastructure etc. This is why CAL does not have a type system of its own, but relies on the environment to provide one (cf. section 4 for more information). We hope that this will allow the use of CAL in a very wide range of contexts, from full-fledged component models (such as JavaBeans) to very resource-constraint embedded platforms.

Making relevant design knowledge explicit and manifest. The key goal of CAL is to enable the author of an actor to express some of the information about the actor and its behavioral properties that are relevant to using the actor (e.g. the verify its appropriate use inside a model, or to generate efficient code from it), but that would be only implicit in a description of the actor in a 'traditional' programming language such as C or Java.

1.3 Platform independence and compatibility

CAL is intended to be adaptable to a variety of different platforms. There are a number of ways to interpret the term 'platform independence', and since this is a very important aspect of the design of CAL, we will discuss our approach to this issue in this section.

For example, it could mean that code written in a language can be run on a

*notions of
platform
independence*

*code
portability*

variety of platforms (which is the interpretation chosen, e.g., in the implementation of the Java programming language). One common approach to achieve code independence would be to define a virtual platform that promises to be implementable in various environments. If this platform is a runtime platform (rather than a source-code level API), this makes not only the source code, but also of the compiled target portable across platforms. This has the obvious advantage of (at least potentially) making every piece of code that is ever written against the virtual platform portable to any implementation, increasing reuse and avoiding fragmentation of the user base. There are, however, downsides to this approach. First, it requires a delicate balance between including all possibly desirable features and selecting those that can be implemented on a wide variety of platforms. If the target platforms are reasonably similar, this may not be a problem, but to the extent that the targets vary, some of them may prevent the inclusion of features that would be very useful on others, resulting in a greatest common denominator design. Second, requiring code and APIs to be independent of any specific platform also makes it harder or impossible to take advantage of platform-specific features.

source vs
target code
portability

Another interpretation of the term focuses on the language and its concepts, rather than the code written in the language. The C language is an example of this: it provides a set of basic concepts, but it leaves many details open to specific implementations, such the sizes and representations of basic data types, and of course the set of library functions used to create programs. As a result, C code itself is not portable, but relies on the presence of a specific set of libraries, and may rely on a specific representation of the data objects. Of course, techniques exist to improve code portability, such as standardization of library sets, and abstraction mechanisms that deal with different data representations. But the general problem with this approach is, of course, that code written in a language is not automatically portable. The advantage, however, is that the language as well as code written in it may be tuned to exploit the specific features of a platform.

language
portability

The design of CAL tries to realize this latter form of portability. The reason is that we intend the language to be used on wide variety of different platforms, and that we do not believe that there is a single abstraction that does justice to all of them—particularly because for some of them, performance of the generated code is a very high priority.

portability in
CAL

Nonetheless, portability of source code as well as target code remains a concern. We intend to improve source code portability by defining profiles for certain classes of platforms, which define things like type systems, basic function and procedure libraries and the like. But these will be based on our experiences with the language in various scenarios, and thus are a second step in the development of the language and its environment. As for target code portability, it seems more reasonable to use existing infrastructure (such as the Java VM, or the Common Language Runtime) wherever possible, rather than developing one from scratch.

platform
profiles

targeting
existing
virtual
platforms

RATIONALE.

Throughout this report, selected aspects of the language design will be discussed in these boxes, presenting the rationale for the decisions made in the specific cases.

These discussions are mainly of interest to language implementors and people interested in the language design. Users only interested in the mechanics of the language may safely ignore them.

IMPLEMENTATION NOTE.

Similarly, discussions of implementation techniques, or aspects of the language that require particular consideration when implementing the language on some platform are set in these boxes.

Part I

Language description

Chapter 2

Introductory remarks

Throughout this part, we will present fragments of CAL syntax along with (informal) descriptions of what these are supposed to mean. In order to avoid ambiguity, we will now introduce a few conventions as well as the fundamental syntactic elements (lexical tokens) of the CAL language.

2.1 Lexical tokens

CAL has the following kinds of lexical tokens:

Keywords. Keywords are special strings that are part of the language syntax and are consequently not available as identifiers. See [B](#) for a list of keywords in CAL.

Identifiers. Identifiers are any sequence of alphabetic characters of either case, digits, the underscore character and the dollar sign that is not a keyword. Sequences of characters that are not legal identifiers may be turned into identifiers by delimiting them with backslash characters.

Identifiers containing the \$-sign are reserved identifiers. They are intended to be used by tools that generate CAL program code and need to produce unique names which do not conflict with names chosen by users of the language. Consequently, users are discouraged from introducing identifiers that contain the \$-sign.

Operators. Operators are written as any string of characters `!, @, #, $, %, ^, &, *, /, +, -, =, <, >, ?, ~, —`. In addition to these strings, a few keywords are used as operators. The set of operators in CAL is extensible, different platforms may augment the set of operators. The language itself defines only a small set of operators—see [appendix C.1](#) for a list of predefined operators.

Delimiters. These are used to indicate the beginning or end of syntactical elements in CAL. The following characters are used as delimiters: (,), {, }, [,], ∴.

Comments. Comments are Java-style, i.e. single-line comments starting with “//” and multi-line comments delimited by “/*” and “*/”.

Numeric literals. CAL provides two kinds of numeric literals: those representing an integral number and those representing a decimal fraction. Their syntax is as follows:¹

```
Integer → DecimalLiteral | HexadecimalLiteral | OctalLiteral
Real → DecimalDigit { DecimalDigit } '.' { DecimalDigit } [Exponent]
      | '.' DecimalDigit { DecimalDigit } [Exponent]
      | DecimalDigit { DecimalDigit } Exponent
DecimalLiteral → NonZeroDecimalDigit { DecimalDigit }
HexadecimalLiteral → '0' ('x'|'X') HexadecimalDigit { HexadecimalDigit }
OctalLiteral → '0' { OctalDigit }
Exponent → ('e'|'E') ['+'|'-'] DecimalDigit { DecimalDigit }
NonZeroDecimalDigit → '1'|'2'|'3'|'4'|'5'|'6'|'7'|'8'|'9'
DecimalDigit → '0'| NonZeroDecimalDigit
OctalDigit → '0'|'1'|'2'|'3'|'4'|'5'|'6'|'7'|'8'
HexadecimalDigit → DecimalDigit
                | 'a'|'b'|'c'|'d'|'e'|'f'
                | 'A'|'B'|'C'|'D'|'E'|'F'
```

RATIONALE.

The reason for allowing identifiers to essentially be any sequence of characters (by providing an ‘escaped’ identifier syntax) is that CAL is intended to interoperate with as many other languages as possible, and therefore cannot assume any particular identifier syntax. We expect that most applications will be using C/C++/Java-style host environments, and thus the lexical conventions of CAL are very similar to those found in these languages. But we did not want to exclude other environments just by a too restrictive choice of the identifier space.

¹In contrast to all other grammar rules in this report, the following rules do not allow whitespace between tokens.

2.2 Typographic conventions

In syntax rules, keywords are shown in **boldface**, while all other literal symbols are enclosed in single quotes.

In examples, CAL code is represented `monospaced` if it is to be read literally. Placeholders inside CAL code, such as for expressions, variables, statements and so forth, are set in italics, as e.g., *E*, *v*, and *S*.

2.3 Syntax rules

We use a form of BNF to describe the syntax rules. Literal elements are put in quotes (in the case of symbols and delimiters), or set in boldface (in the case of keywords). An optional occurrence of a sequence of symbols *A* is written as [*A*], while any number of consecutive occurrences (including none) are written as {*A*}. The alternative occurrence of either *A* or *B* is expressed as *A* | *B*.

We often use plural forms of non-terminal symbols without introducing them explicitly. These are supposed to stand for a comma-separated sequence of at least one instance of the non-terminal. E.g., if *A* is the non-terminal, we might use *As* in some production, and we implicitly assume the following definition:

$$As \rightarrow A \{ \text{'}, ' A \}$$

In the examples we will give in this report, we will assume the 'usual' interpretation of expression literals and mathematical operators, even though strictly speaking these are not part of the language and depend on the environment. A specific implementation of CAL may not have these operators, or interpret them or the literals in a different manner.

2.4 Notational idioms

Like most programming languages, CAL involves a fair number of syntactical constructs that need to be learned and understood by its users in order to use the language productively. The effort involved in gaining familiarity with the language can be a considerable impediment to its adoption, so it makes sense to employ general guidelines for designing the syntax of constructs, which allow users to make guesses about the syntax if they are unsure about the details of a specific language construction. We call these guidelines, which define the style of a language, its notational idioms.

The following is a list of notational idioms guiding the design of CAL's language syntax. Not all of them may hold true all the time, but we tried to conform to them whenever reasonable.

Keyword constructs. Many constructs in CAL are delimited by keywords rather than more symbolic delimiters—such constructs are called keyword constructs. Examples of these would be lambda-expressions (see section 6.9.1). Other constructs are delimited by symbols (e.g. comprehensions, see section 6.10), or are at least partially lacking delimiters (such as assignments, which begin with a variable name, see section 7.1).

Alternative end markers. Every keyword construct ends with either the keyword **end** or the keyword that consists of appending the opening keyword to **end**. For instance, a lambda-expression can either be written

```
lambda (x) : x * x end
```

or alternatively

```
lambda (x) : x * x endlambda
```

Expression head/body separator. Composite expressions often consist of the opening keyword, a head, a separator, the body, and the closing end marker keyword. In such expressions, the separator is usually the ':'-character, as in the following examples:

```
let x = f(k) : g(x, x) end
```

or

```
lambda (x) : x + x end
```

Statement head/body separator. Many statements have a similar structure as the one for expressions. For statements, the keywords **do** or **begin** are used as a separator:

```
while n > 0 do k := f(k); n := n - 1; end
```

or

```
procedure p (x) begin
  println("Result: " + x.toString());
end
```


Chapter 3

Structure of actor descriptions

Each actor description defines a named kind of actor. Actors may refer to entities defined in the implementation context, which provides a hierarchical namespace for these entities, see section 3.1 for details. Actor descriptions may use import declarations to use parts of this namespace or the objects defined in it as part of their global environment.

Actors are the largest lexical units of specification and translation. The basic structure of an actor is this:

```
Actor → [Imports] actor ID
      [ '[' TypePars ']' ] '(' ActorPars ')' IOSig [TimeClause] ':'
      { VarDecl | Action | InitializationAction | PriorityBlock }
      [ActionSchedule]
      { VarDecl | Action | InitializationAction | PriorityBlock }
      (end|endactor)

TypePar → ID [ '<' Type ]
ActorPar → [Type] ID ['=' Expression]
IOSig → [PortDecls] '==>' [PortDecls]
PortDecl → [multi] [Type] ID
TimeClause → time Type
```

The header of an actor contains optional type parameters and actor parameters, and its port signature. This is followed by the body of the actor, containing a sequence of state variable declarations (section 5.1), actions (chapter 8), initialization actions (section 8.5), priority blocks (section 9.3), and at most one action schedule (section 9.2).

Type parameters are variable symbols that are bound to types when the actor is instantiated. They can be used like any other type inside the actor definition. Each type parameter may be optionally bounded, i.e. they may be

type
parameters
type bounds

associated with some type. In this case, the actual type that this parameter is instantiated to is required to be a subtype of the bound (including the bound itself).

By contrast, actor parameters are values, i.e. concrete objects of a certain type (although, of course, this type may be determined by a type parameter). They are bound to identifiers which are visible throughout the actor definition. Conceptually, these are non-assignable and immutable, i.e. they may not be assigned to by an actor.

actor
parameters

IMPLEMENTATION NOTE.

A specific implementation such as the one in Ptolemy might change these parameters, for example in response to user interaction during design. For this to make sense in CAL, the implementation has to ensure the consistency of the actor state with the new parameter values, which it usually does by reinitializing the actor whenever one of its parameters is assigned a new value.

The port signature of an actor specifies the input ports and output ports, including their names, whether the port is a multiport or a single port, and the type of the tokens communicated via the port. While single ports represent exactly one sequence of input or output tokens, multiports are comprised of any number of those sequences (called channels), including zero.

single ports
vs. multiports

3.1 Namespaces and imports

An actor description may contain free variables, i.e. references to variables not defined inside the actor. Often, these are functions or procedures, but also types, which are predefined as part of the respective implementation context. The collection of all globally visible variable bindings is called the global environment of an actor.

However, implementation contexts may be very rich, providing a large number of functions, procedures, and types for actor writers to use. In such cases it would be inappropriate to define all of these as global environment—it would lead to a very large number of variable names, only a very small part of which would actually be used by each actor definition. For this reason, implementation contexts may use a hierarchical namespace for naming these entities, where each entity is denoted by a sequence of identifiers separated by dots (a so-called qualified identifier). Actor specifications may use them as part of their global environment by importing them into it. Effectively, one can think of import declarations as constructing the global environment of an actor description, starting with the default global environment, and adding bindings to it.

The qualified identifiers that denote each entity in the hierarchical namespace have two parts: the (possibly empty) sequence of identifiers up to and excluding

global
environment

hierarchical
context
namespace

imports
assemble
global
environment
qualified id =
subnamespace
+ local name¹⁸

the last, and the last identifier. The first part is called the subnamespace or package, while the second is called the local name. For example, in the qualified identifiers `X.Y.Z`, `XYZ`, and `java.util.HashMap`, the subnamespaces are `X.Y`, `λ`, and `java.util`, respectively,¹ while the corresponding local names are `Z`, `XYZ`, and `HashMap`.

An import declaration can either make a single entity available as the value of a global variable, or the group of all entities inside the same subnamespace.

single vs
group imports

Import \rightarrow SingleImport | GroupImport `;`

SingleImport \rightarrow **import** QualID [`'=` ID]

GroupImport \rightarrow **import all** QualID

QualID \rightarrow ID { `.` ID }

For a single import declaration, the qualified identifier denotes the entity to be imported into the global environment. If the optional identifier following it after an `'=` is omitted, the entity denoted by the qualified identifier is imported under its local name. For instance, the import declaration

```
import A.B.C;
```

imports the entity denoted by `A.B.C` under the name `C` into the global environment. If an identifier is specified, it will be the name of the specified entity:

```
import A.B.C = D;
```

imports the same entity under the name `D`.

Group import declarations import entire groups of entities. In this case, the qualified identifier specifies the subnamespace, and all entities in that subnamespace are imported under their local names. For example,

```
import all A.B;
```

imports `A.B.C` as `C` and `A.B.E` as `E`, if these are the two entities in that subnamespace.

3.2 Time

CAL supports an abstract notion of time as a way to relate the “duration” of various actor firings to each other, and thus potentially control or constrain the concurrency in a model that consists of several actors. It is essential to the division of responsibility between an actor and its environment that the actor itself does not interpret this time information, as it impacts the coordination of more than one actor, rather than the execution of an isolated actor. For this reason, time is purely declarative in CAL, and the language itself does not attach any meaning to it whatsoever, other than it being a property of actor transitions.

time
semantics
external to
actor

In section 10.2 the time system of an actor is introduced as an algebraic

¹ λ denotes the empty sequence of identifiers.

time system:
tags, delays,
addition, zero
delay

structure that has

- a set of time tags
- a set of time delays
- a (partial) order on the tag set
- an addition operation that adds delays to tags to yield larger tags
- a zero delay that is the neutral delay with respect to that addition.

The actor specification, however, only contains the delays. For this reason, the optional **time**-tag in the actor header is only followed by one type, the type of the delay values in action descriptions (see section 8.3). This type must have the following properties:

delay type

- There must be a partial order that is compatible with the addition operation used to add the time delay to tags. More precisely, for any tag t , and any two delays d_1, d_2 , the following must hold:

$$d_1 < d_2 \implies (t + d_1) < (t + d_2)$$

- There must be a zero delay, say z , which defines the set of valid delays of the specified type as follows. If d is of the delay type, it is valid iff $z \leq d$.

The time-clause in the actor head functions as a type declaration for the delays. Its presence does not imply that any or all actions have non-zero delays, and neither does its absence imply that all actions have zero delays—unless a platform requires type information to be added, in which case an actor that contained delay-clauses in actions but no time-clause would not typecheck and hence not be well-formed.

time-clause:
static type
declaration

See section 8.3 on how to specify time delays in actions.

Chapter 4

Data types

CAL is optionally typed, i.e. it allows programmers to give each newly introduced identifier a type (see section 5.1 for more details on declaring variables), but it does not require it, in which case the identifier is said to be untyped. In general, the type system is considered part of the external environment that we try to keep CAL actor specifications orthogonal to. In this chapter we will therefore discuss primarily the syntax provided for writing types, leaving the concrete interpretation to the description of CAL implementations on individual platforms.

optionally
typed

However, CAL does assume a few basic types (some of which are parametric), viz. those that are used as part of some language constructions. There are also some framework rules on how to type checking/inference is to be conducted, which will be discussed in section 4.6.

minimal type
system

4.1 Objects, variables, and types

In general, there are really two kinds of types—the types with which variables are declared (variable types),¹ and the types of runtime objects (object types). In most languages, these are either the same sets, or there is a significant overlap. However, even in common languages, these sets are not identical: in Java, e.g., abstract classes and interfaces are only variable types, never the types of objects.

¹We use *variable* here to mean any name inside the language, including parameters etc.

RATIONALE.

Most programming languages either require explicit typing, or they do not have constructs for expressing types as part of the program source. Some languages perform type inference, i.e. they allow users to omit type declarations as long as they can infer enough information from the context to guarantee the safeness of the program.

Making types in CAL optional reflects our choice to allow for a wide range of embeddings and applications of the language. CAL could thus be used as a scripting language that is interpreted and completely untyped, or as a language that is used to generate C code or hardware from, and which requires complete typing.

Tools are free to reject insufficiently typed programs, or to ignore type annotations if they are not helpful. Of course, tools should make a best effort to infer types whenever they need to know a type that has not been explicitly specified, so as to make actors as reusable as possible across platforms. They should also try to check types whenever feasible, in order to locate errors and to detect malformed actor definitions.

Each variable or parameter in CAL may be declared with a variable type. If it is, then this type remains the same for the variable or parameter in the entire scope of the corresponding declaration. Variable types may be related to each other by a subtype relation, \prec , which is a partial order on the set of all variable types. When for two variable types t, t' we have $t \prec t'$, then we say that t is a subtype of t' , and t' is a supertype of t . Furthermore, a t^2 may be used anywhere a t' can be used—as a consequence, variables of subtypes are substitutable for those of supertypes.

subtype
relation

substitutability

It is important that each object has precisely one object type. As a consequence, object types induce an exhaustive partition on the objects, i.e. for any object type t we can uniquely determine the "objects of type t ".

objects of type
 t

IMPLEMENTATION NOTE.

Stating that each object has an object type does not imply that this type can be determined at run time, i.e. that there is something like run-time type information associated with each object. In many cases, particularly when efficiency is critical, the type of an object is a compile-time construct whose main use is for establishing the notion of assignability, i.e. for checking whether the result of an expression may legally be stored in a variable. In these scenarios, type information is removed from the runtime representation of data objects.

²More precisely, "a t " is an object whose object type is assignable to t , see below.

For each implementation context we assume that there is a set T_V of variable types and T_O of object types. They are related to each other by an assignability relation $\leftarrow_C T_V \times T_O$ which has the following interpretation: for any variable type t_V and object type t_O , $t_V \leftarrow t_O$ iff an object of type t_O is a legal value for a variable of type t_V .

assignability

The assignability relation may or may not be related to subtyping, but at a minimum it must be compatible with subtyping in the following sense. For any two variable types t_V and t'_V , and any object type t_O :

$$t_V \prec t'_V \wedge t'_V \leftarrow t_O \implies t_V \leftarrow t_O$$

subtyping and
assignability

In other words, if an object type is assignable to a variable type, it is also assignable to any of its supertypes.

4.2 Type formats

Even though the CAL language itself does not specify the meaning of most types, it provides notation for expressing types, so actor writers may put this information into their actor descriptions, and an implementation context may use it. There are three basic ways to express types in CAL, and two more constructs for expressing the types of procedural and functional closures (cf. sections 6.9.1 and 6.9.2).

```
Type → ID
| ID '[' TypePars ']'
| ID '(' [TypeAttr { ',' TypeAttr }] ')'
| '[' [Types] '--->' Type ']'
| '[' [Types] '--->' ']'

TypeAttr → ID ':' Type
| ID '=' Expression
```

A type that is just an identifier either refers to a type parameter (if it occurs in the type parameters list of the actor), or it denotes the name of some other non-parametric type. Examples may be `String`, `Integer`.

The form `T[T1, ..., Tn]` is intended to stand for a parametric type T taking the types T_i as parameters. Such a parametric type is also called a type constructor. Many built-in types are of this kind, e.g. `List[Integer]` is a list of elements of type `Integer`, or `Map[String, Real]` is a finite map from keys of type `String` to values of type `Real`.

The next form can be thought of as a more general version of the previous one, where the type constructor has named parameters that may be bound to either types or values. For instance, the type `Matrix[element: Real, width = 4, height = 5]` might be the type of all matrices of real numbers of a certain size.

The type of a lambda closure is written as $[T_1, \dots, T_n \rightarrow T]$, where the T_i are the argument types, and T is the return type. Similarly, the type of a procedural closure is written as $[T_1, \dots, T_n \rightarrow]$, with the T_i again being the argument types.

4.3 Required types

Required types are the types of objects created as the result of special language constructions, usually expressions. The following is a partial list of the required types in CAL, cf. the appendix C for more details on basic data types, their properties, and the operations on them.

- `Null`—the type containing only the value `null`.
- `Boolean`—the truth values `true` and `false`.
- `ChannelID`—the data type comprising the identifiers of channels. Most likely this will be a synonym for some other simple type, like the one of integer or natural numbers.
- `Collection[T]`—a finite collection of elements of type `T`.
- `Seq[T]`—a sequence (finite or infinite) of elements of type `T`.
- `List[T]`—finite lists of elements of type `T`. Lists are subtypes of the corresponding sequences and also subtypes of the corresponding collections, i.e.
`List[T] < Seq[T]` and `List[T] < Collection[T]`
- `Set[T]`—finite sets of elements of type `T`. Sets are subtypes of the corresponding collections, i.e.
`Set[T] < Collection[T]`
- `Map[K, V]`—maps from keys of type `K` to values of type `V`.

pseudotype
`ChannelID`

In addition to these, the types of functional and procedural closures are, of course, also built-in datatypes.

The choice of many fundamental types, such as those representing numeric entities, as well as the operations on them is left to the environment. The only built-in support for these types are the literals which are a result of the lexical scanning of the actor text. Interpretation as well as type assignment to these literals is left up to the environment. See section 6.1 for more details, including a design rationale.

4.4 Structured objects and their types

Many data objects contain other data objects—consider e.g. lists, which contain their elements. An object that contains other objects is called composite or structured, and so is its type.

composite/structured
objects

CAL provides two mechanisms for identifying a subobject inside a structured object (in other words, specifying its location inside the composite): fields and indices. The mechanism used to identify locations inside a composite, as well as the fields and indices that are valid for specifying a location (see below), depend on the composite object and its type.

locations =
fields \cup
indices

The location of a subobject can be used to either retrieve it from the composite (see sections 6.4 and 6.5) or, in the case of mutable objects and types, to replace the subobject with another object (see sections 4.5, 7.1.2, and 7.1.3).

Fields. Fields are names for subobjects inside a composite. Any given object type provides a finite number of those field names, and using any other name to identify a location inside the composite is an error. Since the names are statically provided as part of the program text (as opposed to being computed at runtime) it is possible to check statically whether a field name is permissible for a given object. If A denotes the composite object, and f is the field name, then

fields:
compile-time
locations

$A.f$

$A.f$

selects the subobject in that field. Another consequence of the fact that fields are provided explicitly is that different fields may be associated with different subobject types, i.e. they may be heterogeneous.

heterogeneous
field locations
indices:
runtime
locations

Indices. Indexing into a composite object is a different mechanism for identifying a location inside it. Indices are a number of objects that can be computed at runtime, and which can be thought of as the coordinates of the location inside of the composite. Which and how many indices are valid for a given object is determined by the object. If A denotes the composite object, and E_1, \dots, E_n are n valid indices, then

$A[E_1, \dots, E_n]$

selects the subobject at the specified location. Since indices are computed at runtime, all indexed locations inside a composite object must be assumed to be of the same static type, i.e. they are homogeneous.

homogeneous
indexed
locations

Example 1. A typical application of fields would be a type `Complex` with fields `real` and `imaginary`.

Lists in CAL are indexed by the natural numbers, starting at 0. For instance, the list `[11, 7, 19]` can be indexed with the indices 0, 1, and 2, so that e.g.

`[11, 7, 19][1]`

would yield the value 7. See section C.2.4 for details.

While natural numbers are very common indices, indexed composites are in general not restricted to numeric indices, but can allow any kind of object for the purpose. See section C.2.5 for details on how maps are indexed into.

indices of any
type

Example 2. A `Map[K, V]` from a key type `K` to a value type `V` accepts as indices objects of type `K`, and indexing is simply the application of the map. Say the map

```
map {"abc" -> 15, "def" -> 7}
```

(of type `Map[String, Integer]`) can be indexed by a string as follows:

```
map {"abc" -> 15, "def" -> 7}["abc"]
```

resulting in the value 15.

The following CAL types are required to support indexing:

- `Seq[T]`—index type is the natural numbers, starting from 0.
- `Map[K, V]`—index type is `K`.

Abstractly, indexing is the invocation of a special function, the indexer of a given datatype. The type of the resulting expression may depend on the number and types of the indices, and of course also on the indexed object.

4.5 Mutable objects and their types

Some structured types allow the modification or mutation of an object, i.e. changing the object at some location inside it. Such a type is called mutable, and so is an object of that type.

Mutating objects without any restrictions would be a technique that would render a program essentially unanalyzable in the general case. For this reason, CAL imposes a number of constraints on the ability to use this feature, which will be discussed in the context of mutable variables in sections 7.1.2 and 7.1.3.

However, CAL makes no guarantee that the various locations of an object are independent, i.e. that all other locations remain unaffected by an assignment. A structured mutable type whose locations are independent is called *free*, its locations are called *orthogonal*—all predefined types in CAL are of this kind if they support mutation.

orthogonal
locations
free types

Example 3. A common non-orthogonal type might be `Complex`, the type of complex numbers. Let us assume its objects are mutable, and each has four locations, represented as fields: `real`, `imag`, `modulus`, and `argument`. If `c` is a mutable variable containing a complex number, then

```
c.real := 4;
```

```
c.imag := 3;
```

would at the same time result in the `modulus` component to be set to 5, and the `argument` to be set to $\arctan 4/3$. Now the assignment

```
c.argument := 2;
```

can be interpreted as leaving the `modulus` unchanged, and modifying the `real` and `imag` fields accordingly, viz. to $5 \cos 2$ and $5 \sin 2$, respectively.

4.6 Type framework

Section 4.1 introduced two relations between types: the subtyping relation \prec between variable types, and the assignability relation \leftarrow between variable types and object types. It also required that the two relations be consistent with each other in the sense that assignability to a variable type would imply assignability to any of its supertypes.

This section introduces some simple properties that any type system provided by an implementation context must have. These properties, together with the basic minimal type system presented in section 4.3, constitute the CAL type framework.

Subtype ordering. The subtyping relation \prec on the variable types is a partial order, i.e. for any $t_1, t_2, t_3 \in T_V$ the following holds:

1. Reflexivity: $t_1 \prec t_1$
2. Anti-symmetry: $t_1 \prec t_2 \wedge t_2 \prec t_1 \implies t_1 = t_2$
3. Transitivity: $t_1 \prec t_2 \wedge t_2 \prec t_3 \implies t_1 \prec t_3$.

Finite cover. For any variable type t , there is a unique finite, and possibly empty, set \mathbf{Supt} of variable types which has the following properties:

1. $t \notin \mathbf{Supt}$
2. $\forall t' \in \mathbf{Supt} : t \prec t'$
3. $\forall t', t'' \in \mathbf{Supt} : t' \not\prec t''$
4. $\forall t' : t \prec t' \implies \exists t'' \in \mathbf{Supt} : t'' \prec t'$

finite set of
direct
supertypes

The second condition ensures that the elements of \mathbf{Supt} are all superclasses of t . The third condition requires the set to be minimal. The third condition requires it to be exhaustive: all proper supertypes of t are supertypes of at least one element in \mathbf{Supt} .

Least upper bound. For any two variable types $t_1, t_2 \in T_V$ there is a unique finite, and possibly empty, set of variable types $t_1 \sqcup t_2$ which has the following properties:

1. $\forall t \in t_1 \sqcup t_2 : t_1 \prec t \wedge t_2 \prec t$
2. $\forall t', t'' \in t_1 \sqcup t_2 : t' \not\prec t''$
3. $\forall t : (t_1 \prec t \wedge t_2 \prec t) \implies \exists t' \in t_1 \sqcup t_2 : t' \prec t$

finite set of
supertypes for
any pair of
types

Again, the first condition ensures that all types in the set are indeed upper bounds of the two types in question. The second condition requires them to be minimal. The third condition requires the set to be exhaustive: all supertypes of t_1 and t_2 that are not in that set are supertypes of a type that is.

Chapter 5

Variables

Variables are placeholders for other values. They are said to be bound to the value that they stand for. The association between a variable and its value is called a binding.

bindings

CAL distinguishes between different kinds of bindings, depending on whether they can be assigned to (assignable variable binding), and whether the object they refer to may be mutated (mutable variable binding—cf. sections 4.5 and 7.1.3).

assignable &
mutable
bindings

This chapter first explains how variables are declared inside CAL source code. It then proceeds to discuss the scoping rules of the language, which govern the visibility of variables and also constrain the kinds of declarations that are legal in CAL.

5.1 Variable declarations

Each variable (with the exception of those predefined by the platform) needs to be explicitly introduced before it can be used—it needs to be declared or imported (section 3.1). A declaration determines the kind of binding associated with the variable it declares, and potentially also its (variable) type. There are the following kinds of variable declarations:

declarations

- explicit variable declarations (section 5.1.1,
- actor parameters (chapter 3),
- input patterns (section 8.1),
- parameters of a procedural or functional closure (section 6.9).

Variables declared as actor parameters, in input patterns, or as parameters of a procedural or functional closure are neither assignable nor mutable.

The properties of a variable introduced by an explicit variable declaration depend on the form of that declaration.

5.1.1 Explicit variable declarations

Syntactically, an explicit variable declaration¹ looks as follows:

$$\text{VarDecl} \rightarrow [\mathbf{mutable}] [\text{Type}] \text{ID} [(\text{'='} \mid \text{' :='}) \text{Expression}] \\ \mid \text{FunDecl} \mid \text{ProcDecl}$$

We will discuss function and procedure declarations (FunDecl and ProcDecl) in section 6.9.3.

An explicit variable declaration can take one of the following forms, where T is a type, v an identifier that is the variable name, and E an expression of type T :

- $T \ v$ —declares an assignable, non-mutable variable of type T with the default value for that type as its initial value. It is an error for the type not to have a default value.
- $T \ v \ := \ E$ —declares an assignable, non-mutable variable of type T with the value of E as its initial value.
- `mutable` $T \ v \ := \ E$ —declares an assignable and mutable variable of type T with the value of E as its initial value.
- `mutable` $T \ v \ = \ E$ —declares a non-assignable and mutable variable of type T with the value of E as its initial value.
- $T \ v \ = \ E$ —declares a non-assignable, non-mutable variable of type T with the value of E as its initial value.

Variables declared in any of the first four ways are called *stateful variables*, because they or the object they are containing may be changed by the execution of a statement. Variables declared in the last way are referred to as *stateless variables*.

stateful vs
stateless

Explicit variable declarations may occur in the following places:

- actor state variables
- the **var** block of a surrounding lexical context
- variables introduced by a **let**-block

While actor state variables and variables introduced in a **var**-block can be state variables as well as non-state variables, a **let**-block may only introduce non-state variables.

¹These declarations are called “explicit” to distinguish them from more “implicit” variable declarations that occur, e.g., in generators or input patterns.

5.2 Variable scoping

The scope of a variable is the lexical construct that introduces it—all expressions and assignments using its name inside this construct will refer to that variable binding, unless they occur inside some other construct that introduces a variable of the same name, in which case the inner variable shadows the outer one.

lexical scoping

In particular, this includes the initialization expressions that are used to compute the initial values of the variables themselves. Consider e.g. the following group of variable declarations inside the same construct, i.e. with the same scope:

initialization expression

```
n = 1 + k,
k = 6,
m = k * n
```

This set of declarations (of, in this case, non-mutable, non-assignable variables, although this does not have a bearing on the rules for initialization expression dependency) would lead to k being set to 6, n to 7, and m to 42. Initialization expressions may not depend on each other in a circular manner—e.g., the following list of variable declarations would not be well-formed:

well-formed declaration

```
n = 1 + k,
k = m - 36,
m = k * n
```

More precisely, a variable may not be in its own dependency set. Intuitively, this set contains all variables that need to be known in order to compute the initialization expression. These are usually the free variables of the expression itself, plus any free variables used to compute them and so on—e.g., in the last example, k depended on m , because m is free in $m - 36$, and since m in turn depends on k and n , and n on k , the dependency set of k is $\{m, k, n\}$, which does contain k itself and is therefore an error.

dependency set

This would suggest defining the dependency set as the transitive closure of the free variable dependency relation—which would be a much too strong criterion. Consider e.g. the following declaration:

```
f = lambda (n) :
  if n = 0 then 1 else n * f(n - 1) end
end
```

Here, f occurs free in the initialization expression of f , which is clearly a circular dependency. Nevertheless, the above definition simply describes a recursive function, and should thus be admissible.

recursion

The reason why f may occur free in its own definition without causing a problem is that it occurs inside a closure—the value of f need not be known in order to construct the closure, as long as it becomes known before we use it—i.e. before we actually apply the closure to some argument.

We will now define the dependency sets I_v and D_v of a variable v among a set of variables V that are defined simultaneously in the same scope.

Definition 1 (I_v, D_v —the dependency sets of a variable v). Consider a set V of variables v which are defined simultaneously, i.e. the initial value of each of these variables defined by an expression E_v which is in the scope of all the variables in V . Let F_v be the set of free variables of E_v . As we are only interested in the free variables in V , we will usually use the intersection $F_v \cap V$.

The *dependency set* D_v is defined as the smallest set such that the following holds:

- (a) $F_v \cap V \subseteq D_v$
- (b) $\bigcup_{x \in D_v} D_x \subseteq D_v$

The *immediate dependency set* I_v of each variable v is defined as follows

$$I_v = \begin{cases} \emptyset & \text{for } E_v \text{ a closure} \\ D_v & \text{otherwise} \end{cases}$$

immediate
dependency
set

Intuitively, D_v contains those variables in V on which the object bound to v directly or indirectly depends. I_v is the set of variables whose values need to be known when the object computed by E_v is created—for most expressions, it is the same as D_v , but for closures (procedural or functional) this set is empty, because there is no need to evaluate the body of the closure in order to construct the closure object.²

Now we capture the notion of well-formedness of a set of simultaneously defined variables V as a condition on the dependency sets as follows:

Definition 2 (Well-formed declaration set). A set of simultaneously declared variables (a *declaration set*) V is *well-formed* iff for all $v \in V$

$$v \notin I_v$$

well-formed
declaration set

Note that, as in the example above, a variable may occur free in its own initialization expression, but still not be in its own immediate dependency set, as this only includes those variables whose value must be known in order to compute the value of the declared variable.

This notion of well-formedness is useful because of the following property:

Corollary 1 (No mutual dependencies in well-formed variable sets.). Given a well-formed variable set V , for any two variables $v_1, v_2 \in V$, we have the following property:

$$\neg(v_1 \in I_{v_2} \wedge v_2 \in I_{v_1})$$

That is, no two variables ever mutually immediately depend on each other.

²Here we use the fact that closures can be constructed without the values of their free variables, which is clearly an artifact of the way we envision closures to be realized, but it is a useful one.

The proof of this property is trivial, by contradiction and induction over the definition of the dependency set (showing that mutual dependency would entail self-dependency, and thus contradict well-formedness).

RATIONALE.

Strictly speaking, this definition of well-formedness is a conservative approximation to what might be considered well-formed variable declarations. It is chosen so that (a) it allows commonly occurring circular dependencies to be expressed and (b) it can be implemented without undue efforts. But it is important to realize that it does declare some declaration sets as erroneous which could easily be interpreted. For instance, consider the following case:

```
f = lambda (x) :
  if x <= 1 then 1 else x * g(x - 1) end
end,
g = f
```

According to the definitions above, both `g` ends up in its own immediate dependency set, and thus these declarations are not well-formed.

This allows us to construct the following relation over a set of variables:

Definition 3 (Dependency relation). Given a set of variables V defined in the same scope, we define a relation \prec on V as follows:

dependency
order relation

$$v_1 \prec v_2 \iff v_1 \in I_{v_2}$$

In other words, a variable is ‘smaller’ than another according to this relation iff it occurs in its dependency set, i.e. iff it has to be defined before the other can be defined. The well-formedness of the declaration set implies that this relation is a non-reflexive partial order, since variables may not mutually depend on each other.

This order allows us to execute variable declarations in such a way that immediate dependencies are always evaluated before the dependent variable is initialized.

Example 4. Consider the following variable definitions:

```
a = f(x),
f = lambda (v) :
  if p(v) then v else g(h(v)) end
end,
```



```
g = lambda (w) : b * f(w) end,
b = k
```

Note that f and g are mutually recursive.

The following lists the immediate dependencies and the free variable dependencies of each variable above,³ along with their intersection with the set $\{a, f, g, b\}$, which is the set V in this case:

v	F_v	$F_v \cap V$	D_v	I_v
a	$\{f, x\}$	$\{f\}$	$\{f, g, b\}$	$\{f, g, b\}$
f	$\{p, g, h\}$	$\{g\}$	$\{f, g, b\}$	\emptyset
g	$\{b, f\}$	$\{b, f\}$	$\{f, g, b\}$	\emptyset
b	$\{k\}$	\emptyset	\emptyset	\emptyset

Now let us compute the dependency set D_a of the variable a . We start with the set

$$F_a \cap V = \{f\}$$

Now we compute

$$(F_a \cup F_f) = \{f, g\}$$

Then

$$(F_a \cup F_f \cup F_g) \cap V = \{f, g, b\}$$

Finally, we reach a fixpoint at

$$D_a = (F_a \cup F_f \cup F_g \cup F_b) \cap V = \{f, g, b\}$$

Similarly, we compute D_f , D_g , and D_b . The immediate dependency sets of f and g are empty, because their initialization expressions are closures. The immediate dependency sets of a and b are the same as their dependency sets, which in the case of b is also empty, because it does not depend on any variable in V . As a result of this analysis, we see that the variables f , g , and b may be defined in any order, but all must be defined before a , as it depends on all of them.

Example 5. Now consider the following slightly changed variable definitions, with an additional dependency added to b :

```
a = f(x),
f = lambda (v) :
  if p(v) then v else g(h(v)) end
end,
```

³We are disregarding here the implicit variable references that will be introduced when the operators are resolved to function calls—strictly speaking, they would become part of F_v , but as they are always referring to global variables, and would thus disappear from $F_v \cap V$ anyway, we do not bother with them in the example.

```

g = lambda (w) :
    b * f(w)
end,
b = a * k

```

Again, the following table lists the dependency sets:

v	F_v	$F_v \cap V$	D_v	I_v
a	$\{f, x\}$	$\{f\}$	$\{f, g, b, a\}$	$\{f, g, b, a\}$
f	$\{p, g, h\}$	$\{g\}$	$\{f, g, b, a\}$	\emptyset
g	$\{b, f\}$	$\{b, f\}$	$\{f, g, b, a\}$	\emptyset
b	$\{a, k\}$	$\{a\}$	$\{f, g, b, a\}$	$\{f, g, b, a\}$

Now, computing D_a proceeds as follows:

$$F_a \cap V = \{f\}$$

$$(F_a \cup F_f) \cap V = \{f, g\}$$

$$(F_a \cup F_f \cup F_g) \cap V = \{f, g, b\}$$

$$(F_a \cup F_f \cup F_g \cup F_b) \cap V = \{f, g, b, a\}$$

$$D_a = (F_a \cup F_f \cup F_g \cup F_b \cup F_a) \cap V = \{f, g, b, a\}$$

Obviously, in this case $a \in I_a$, thus the set of variable definitions is not well-formed.

Chapter 6

Expressions

Expressions evaluate to a value and are side-effect-free, i.e. they do not change the state of the actor or assign or modify any other variable. Thus, the meaning of an expression can be described by the value it is evaluating to.

no side effects

If the computation of an expression terminates, it results in a value, and that value has an object type. The precise value depends on the environment surrounding the expression when it is evaluated. In general, the objects computed by an expression in different environments may belong to different object types. However, Using the properties of the type system outlined in section 4.6, we can compute a set of least upper variable type bounds for each expression, knowing the declared types of each of its free identifiers and the types of the literals occurring in it. The object types of all objects ever computed by the expression must be assignable to all of the variable types in that set.

expressions
and types

The following is an overview of the kinds of expressions and expression syntaxes provided in CAL.

expression
syntax

Expression \rightarrow PrimaryExpression { Operator PrimaryExpression }

PrimaryExpression \rightarrow [Operator] SingleExpression
{ '(' [Expressions] ')' | '[' Expressions ']' | '.' ID }

SingleExpression \rightarrow [old] ID
| ExpressionLiteral
| '(' Expressions ')'
| IfExpression
| LambdaExpression
| ProcExpression
| LetExpression
| ListComprehension | SetComprehension | MapComprehension
| TypeAssertionExpr

We will now discuss the individual kinds of expressions in more detail.

6.1 Literals

Expression literals are constants of various types in the language. They look as follows:

literals
describe
constants

```
ExpressionLiteral → IntegerLiteral | DecimalFractionLiteral
                  | StringLiteral
                  | true | false | null
```

The type of **true** and **false** is `Boolean`, the type of **null** is `Null`.

The exact types of the other literals are determined by the environment, but the intuition behind them is that they represent the integers, the real numbers, and character strings, respectively.

literal types
determined by
context

RATIONALE.

We are not committing to particular types for the numeric literals because we want to keep the requirements on the type system as weak as possible, to allow for a wide range of implementations. For example, some implementations may have only integer numbers up to a specific word size (say, 32 or 64 bits), while others have variable-sized integers. Such an implementation may want to assign a specific type to an integer literal depending on its numeric size—e.g., it may make it an object of type `Integer32` if it fits into 32 bits, and of type `Integer` if it is larger than that.

6.2 Variable references

The expression used to refer to the value bound to a variable at any given point during the execution is simply the name of the variable itself, i.e. an identifier.

6.2.1 Old variable references

The code inside an action may refer to the value of a variable at the beginning of the action by prefixing the variable name with the keyword **old**, as in the following example:

assignable &
mutable
bindings

Example 6. . . .

```

sum := 0;

action [a] ==> [old sum / sum]
do
  sum := sum + a;
end

...

```

The output expression refers to both the value of `sum` at the beginning of the firing as well as its value at the end of the firing. This code is equivalent to the following:

```

...

sum := 0;

action [a] ==> [oldSum / sum]
var
  oldSum = sum
do
  sum := sum + a;
end

...

```

old-
references:
state *before* a
firing

In other words, using **old** values of variables in an action introduces an implicit non-assignable, non-mutable variable (cf. section 7.1.4 for the implications of the original variable being mutable).

implicit
variable

Closures created inside an action may also refer to **old** variables, and the meaning of this follows from the transformation above: They will always refer to the value of the variable at the beginning of the firing that created them.

old-references
and closures

The **old** keyword may not be used outside of an action, or in front of a variable that is not an assignable or mutable actor state variable.

6.3 Function application

Function applications have the following syntactic format:

FunctionApplication \rightarrow Expression '(' Expressions ')'

The first expression must evaluate to a function of the appropriate type and arity. The expression inside the parentheses are the function arguments. The value of the function application expression is the value computed by the function when applied to the values of the argument expressions.

language imple-
mentation: the
art of
cheating 37
without
getting caught

IMPLEMENTATION NOTE.

Functions may, in principle, be non-strict, i.e. evaluation of argument expressions may be deferred until their values are actually needed. The bounds for implementations are defined only by an expression's proper behavior with respect to the modification of state. For example, consider the following statements:

```
g := lambda (x) : lambda (y) : x + y end end;
a := 11;
f := g (a * a);
a := 7;
b := f(1);
```

Any correct implementation must ensure that `b` has the value 122 at the end of this statement sequence. If it naively defers the evaluation of the argument expression `a * a` until `f` is called, `b` will erroneously be assigned the value 50. This, of course, does not mean that the argument expression must be evaluated when `g` is called, but it must behave as if it were evaluated in the state valid at the time of the call whenever its value is needed. In this example, it could move the old value of `a` into the lambda-closure and use it, rather than the current value of `a`, when it finally evaluates the argument expression.

In other words, if an implementation chooses to employ some form of lazy evaluation, the effect is only to give meaning to programs that would otherwise not be valid, while not affecting the meaning of programs that do evaluate under an eager implementation. Consider this example:

```
let
  f = lambda (x, y) : x end,
  g = lambda () : g() end
:
  f(1, g())
end
```

In lazy implementations, this expression will yield 1, while in eager implementations it will be an error.

6.4 Field selection

A field selector expression extracts a sobobject from a composite object (see 4.4 for details on composite objects). The syntax is as follows:

FieldSelectorExpr \rightarrow Expression '.' ID

The result of this expression is the subobject of the value of the expression that is contained in the field specified by the identifier.

6.5 Indexing

An indexing expression selects a subobject from a composite object (cf. section 4.4 for more details). Syntactically, indexing expressions are similar to function applications, although they use square brackets for enclosing the arguments. The general format is

indices define
location

$$\text{IndexerExpr} \rightarrow \text{Expression } '[' \text{ Expressions } ']$$

where the first expression must be of a type that supports an indexer, and the expressions between the brackets must be indices specifying a valid location for the given object. The type of an indexing expression is determined by the indexer, which is different for each structured data type, and may differ according to the number of indices and their types.

6.6 Operators

There are two kinds of operators in CAL: unary prefix operators and binary infix operators. A binary operator is characterized by its associativity and its precedence. In CAL, all binary operators associate to the left, while their precedence is defined by the platform. Unary operators always take precedence over binary operators.

unary &
binary
operators

Example 7.

$a + b + c$ is always $(a + b) + c$.

$\#a + b$ is always $(\#a) + b$.

$a + b * c$ is $a + (b * c)$ if $*$ has a higher precedence than $+$, which is usually the case (see also the appendix C.1).

Operators are just syntactical elements—they represent ordinary unary or binary functions, so the only special rules for operators are syntactical. In general, the set of operators is defined by the implementation context, although a small number of operators are predefined. These operators are represented by keywords, as opposed to strings of symbols, which represent all the other operators (cf. section 2.1).

syntactical
sugar

6.7 Conditional expressions

The simple conditional expression has the following form:

IfExpression \rightarrow **if** Expression **then** Expression **else** Expression **end**

The first subexpression must be of type `Boolean`, and the value of the entire expression is the value of the second subterm if the first evaluated to `true`, and the value of the third subterm otherwise.

The type of the conditional expression is the most specific supertype (least upper bound) of both, the second and the third subexpression. It is undefined (i.e. an error) if this does not exist.

6.8 Introducing a local scope

In expressions, local variables are introduced using a **let**-construct. This is often useful to factor out large subexpressions that occur several times.

LetExpression \rightarrow **let** VarDecls ':' Expression (**end**|**endlet**)

The list of local definitions defines new identifiers and binds them to values. The variables (which are non-mutable and non-assignable) are visible inside the body expression. Its type is the type of the entire construct.

6.9 Closures

Closures are objects that encapsulate some program code along with the variable context (its environment) that was valid when it was created. CAL distinguishes two kinds of closures, which differ in the kind of code they encapsulate:

closure = code
+
environment

- function closures (or simply functions) contain a parametric expression,
- procedure closures (or just procedures) contain a parametric list of statements.

functions &
procedures

The two kinds of closures are used in different contexts, and in different ways—the application of a functional closure to (a tuple of) arguments is an expression (cf. section 6.3), whereas the call of a procedural closure on (a tuple of) arguments is a statement (cf. section 7.2).

6.9.1 Lambda-expressions and function closures

Function closures are the result of evaluating a **lambda**-expression. They represent functions that are defined by some expression which is parameterized and may also refer to variables defined in the surrounding context.

function =
expression +
environment

$\text{LambdaExpression} \rightarrow [\mathbf{const}] \mathbf{lambda} \text{'(' [FormalPars] \text{'})' ['-->' Type]}$
 $\quad [\mathbf{var} \text{ VarDecls} \text{' ':' Expression (end|endlambda)}$
 $\text{FormalPar} \rightarrow [\text{Type}] \text{ID}$

Applying function closures is side-effect free, i.e. their application (to arguments) does not change the state. However, in general they may refer to stateful variables, and thus may themselves depend on the assignment of variables in their context, and thus be affected by side effects of other constructs.

function
application
side-effect-
free
invariant
function
closures

The **const** keyword identifies those function closures for which this is not the case, i.e. which do not refer to variables whose values may change—such a function is also called an invariant function (closure). It does not change the behavior of the closure, i.e. removing it will not affect the value computed by the closure. It is intended to serve as a declaration that expresses the programmers intention, and that may be checked by a compiler. It is an error for a **const lambda**-closure to refer to assignable or mutable variables.

If the types of the formal parameters are T_1 to T_n , respectively, and the return type is T , then the type of the lambda expression is

$[T_1, \dots, T_n \text{ --> } T]$

The type of an invariant function closure is a subtype of the corresponding function closure type, and is written as

$\mathbf{const} [T_1, \dots, T_n \text{ --> } T]$

The only built-in operation defined on a function closure is its application to a tuple of arguments, cf. section 6.3.

applying
functions \rightarrow
6.3

6.9.2 Proc-expressions and procedure closures

Procedure closures are somewhat similar to function closures, in that they encapsulate a piece of code together with the context in which it was defined. However, in the case of procedure closures, this piece of code is a list of statements, i.e. executing a procedure closure is likely to have side effects (as opposed to the application of a function closure).

procedure =
statements +
environment

Syntactically, a procedure closure looks as follows:

$\text{ProcExpression} \rightarrow \mathbf{proc} \text{'(' [FormalPars] \text{'})' [\mathbf{var} \text{ VarDecls}]}$
 $\quad (\mathbf{do} | \mathbf{begin}) \{ \text{Statement} \} (\mathbf{end} | \mathbf{endproc})$

If the types of the formal parameters are T_1 to T_n , respectively, then the type of the proc expression is

$[T_1, \dots, T_n \text{ -->}]$

Since block closures can produce side effects, their execution cannot be part

calling
procedures
 \rightarrow 7.2

of the evaluation of an expression. Executing a block closure is a fundamental kind of statement, which is discussed in section 7.2.

6.9.3 Function and procedure declarations

One very common use for closures is the definition of functions or procedures with a particular fixed name inside some scope, often the actor itself. This can be done using the standard variable declaration syntax, as follows:

```
timestwo = lambda(x) : 2 * x end
```

However, since this use is so frequent, CAL provides special syntax that looks a little more familiar and makes the definition of functions and procedures a little easier. The above could also be written like this:

```
function timestwo (x) : 2 * x end
```

The general format for these constructs is as follows:

```
FuncDecl → function ID '(' [FormalPars] ')' [var VarDecls ':']  
          Expression end
```

```
ProcDecl → procedure ID '(' [FormalPars] ')' [var VarDecls (begin | do)]  
          { Statement } end
```

The variable introduced by these declarations is non-assignable and non-mutable.

6.10 Comprehensions

Comprehensions are expressions which construct one of the built-in composite objects: sets, lists, or maps. There are two variants of comprehensions, those with and those without generators. We will first focus on the simpler ones without generators, and then turn to the more general comprehensions with generators. The reason for this order of presentation is that the meaning of comprehensions with generators will be defined by reducing them to simple collection expressions.

constructing
composite
objects

Note. The `Collection` type is a supertype of both `Set` and `List`, but not `Map`. In spite of this, we will use the term `collection` for maps as well in this section, because the way they are constructed is very much the same as for sets and lists. To avoid confusion, we will refer to sets and lists as proper collections or `Collections` if we want to distinguish them from maps.

6.10.1 Simple collection expressions

Simple collection expressions just enumerate the elements of the set or list and the mappings of the map, respectively. They are written as follows:

simple comprehension:
enumerate elements

SimpleSetComprehension \rightarrow `'{ [Expressions] }'`

SimpleListComprehension \rightarrow `'[[Expressions ['|' Expression]]]'`

SimpleMapComprehension \rightarrow `map '{ [Mappings] }'`

The elements of a set are not ordered, and each element occurs inside the set at most once. If two or more element expressions evaluate to the same value, they result in only one element inside the set.¹ In contrast, the elements in a list are ordered, and the same element may occur in it more than once. Also, the list syntax allows a tail to be specified, which must be a list of the same type that is appended to the list of the explicitly specified elements.

unordered set
vs ordered list

Example 8. If n is the number 10, then the simple set expression

`{n, n*n, n-5, n/2}`

evaluates to the set `{10, 100, 5}`.

If s represents the list `[1, 2, 3]`, then the simple list expression

`[4, 5, 6 | s]`

is the list `[4, 5, 6, 1, 2, 3]`.

Simple map expressions explicitly specify the mappings from keys to values. Similar to the case of sets, if two key expressions result in the same value, only one key/value mapping will be generated from them. If the corresponding value expressions are not the same, one of the values will be chosen.

nondeterminism
in map
construction

Example 9. Let again n be 10. The map

`map {n -> 1, n*n -> 2, n-5 -> 3, n/2 -> 4}`

evaluates to either `map {10 -> 1, 100 -> 2, 5 -> 3}`

or to `map {10 -> 1, 100 -> 2, 5 -> 4}`.

6.10.2 Comprehensions with generators

Simple comprehension expressions only allow the construction of sets, lists, or maps of a size that is directly correlated with the size of the expression. In order to facilitate the construction of large or variable-sized collections, CAL provides generators to be used inside an expression constructing a collection. The syntax of the resulting general form of comprehensions looks as follows:

constructing
large
collections

¹For this reason, it is in general only possible to infer an upper bound on the number of elements from a simple set expression, not the precise number of elements.

```

SetComprehension → '{' [Expressions ':' Generators] '}'
ListComprehension → '[' [Expressions ':' Generators] '[' Expression ']' ']'
MapComprehension → map '{' [Mappings ':' Generators] '}'
Mapping → Expression '->' Expression
Generator → for [Type] ID '[' IDs] in Expression { ',' Expression }

```

The generators, which begin with the `for` keyword, introduce new variables, and successively instantiate them with the elements of the proper collection after the `in` keyword. The expression computing that collection may refer to the generator variables defined to the left of the generator it belongs to. If that expression is of type `Collection[T]`, the corresponding generator variable is of type `T` unless declared otherwise.

generators &
filters

The optional expressions following the collection expression in a generator are called filters—they must be of type `Boolean`, and only variable bindings for which these expressions evaluate to **true** are used to construct the collection.

Example 10. The expression `{}` denotes the empty set, while

```
{1, 2, 3}
```

is the set of the first three natural numbers. The set

```
{2 * a : for a in {1, 2, 3}}
```

contains the values 2, 4, and 6, while the set

```
{a : for a in {1, 2, 3}, a > 1}
```

describes (somewhat redundantly) the set containing 2 and 3. Finally, the set

```
{a * b : for a in {1, 2, 3},
         for b in {4, 5, 6}, b > 2 * a}
```

contains the elements 4, 5, 6, 10, and 12.

Writing the above as

```
{a * b : for a in {1, 2, 3},
         b > 2 * a, for b in {4, 5, 6}}
```

is **illegal** (unless `b` is a defined variable in the context of this expression, in which case it is merely very confusing!), because the filter expression

```
b > 2 * a
```

occurs before the generator that introduces `b`.

generator
variable
scoping

If the generator collection is a set rather than a list, the order in which elements are extracted from it will be unspecified. This may affect the result in case of a list comprehension.

element order
in generators

Example 11. Because lists are order-sensitive, the list

```
[ a : for a in [1, 2, 3] ]
```

is different from the list

```
[ a : for a in [3, 2, 1] ]
```

If the collection computed in a generator is *not* itself a list but a set, as in

[a : for a in {1, 2, 3}]
 then the order of the elements in the resulting list will be indeterminate.

Generator semantics. In order to precisely describe the evaluation of a comprehension that contains generators and filters, we need to introduce a few symbols first. We call the original expression E . Removing all generators from E results in the simple collection expression E' . Now E has $N \geq 0$ generators. The i -th generator has $M_i \geq 1$ variables, with names $v_{i,1}, \dots, v_{i,M_i}$. The collection expression of the i -th generator will be called C_i . Following the i -th generator are $K_i \geq 0$ filter expressions, which we call $F_{i,1}, \dots, F_{i,K_i}$. Let Z denote the corresponding empty collection expression, i.e. $\{\}$ for sets, $[\]$ for lists, and `map { }` for maps. The tail expression of lists will be treated separately: we replace $[S|T]$ with $[S] + T$ first, and then apply the following algorithm to the tail-free list comprehension $[S]$.

for
 $v_{i,1}, \dots, v_{i,M_i}$ in
 C_i ,
 $F_{i,1}, \dots, F_{i,K_i}$

In the following, we define the meaning of generators in comprehensions by replacing them with a number of previously described constructs, such as function closures, function application, and conditional expressions. The key to this replacement is a function `$mapadd`, which we will describe below.

$$\begin{aligned}
 E &\equiv G(1) \\
 G(i) &\equiv \begin{cases} E' & i > N \\ \text{let } a_i \equiv C_i \text{ ; } G_V(i, 1) \text{ end} & \text{otherwise} \end{cases} \\
 G_V(i, j) &\equiv \begin{cases} G_F(i, 1) & j > M_i \\ \$\text{mapadd}(a_i, \text{lambda } (v_{i,j}) : G_V(i, j + 1) \text{ end}) & \text{otherwise} \end{cases} \\
 G_F(i, k) &\equiv \begin{cases} G(i + 1) & k > K_i \\ \text{if } F_{i,k} \text{ then } G_F(i, k + 1) \text{ else } Z \text{ end} & \text{otherwise} \end{cases}
 \end{aligned}$$

syntactical
 transforma-
 tion of
 generators

The a_i are distinct and fresh variable symbols, i.e. they are mutually different, and they also differ from any variable symbol occurring in the comprehension.

The `$mapadd` function takes two arguments, a collection and a unary function. It iterates over the elements of its first argument, and applies the function to each element. The resulting values are added—for sets, this means it produces the union of all results, for lists the concatenation, and for maps the map-union.

`$mapadd:`
`Collection[T]`
`[T-->Set/List/Map[T]]`

RATIONALE.

Of course, whether there exists a function `$mapadd` at all, and whether it has this name will be implementation-dependent. The reason for choosing a name with a `$`-sign is because users are discouraged from using the `$`-sign for their identifiers, so that generated identifiers and internal names can be formed without creating conflicts with names chosen by the user.

Example 12. The expression

```
$mapadd([1, 2, 3], lambda(x) : {x, x + 2} end)
results in {1, 3} + {2, 4} + {3, 5}, which is
{1, 2, 3, 4, 5}
```

The expression

```
$mapadd([1, 2, 3], lambda(x) : [x, x + 2] end)
results in [1, 3] + [2, 4] + [3, 5], which is
[1, 3, 2, 4, 3, 5]
```

The expression

```
$mapadd([1, 2, 3],
lambda(x) : map{x->x*x, x+2->x} end)
results in map{1->1, 3->1} + map{2->4, 4->2} + {3->9, 5->3}, which
is either
map{1->1, 2->4, 3->1, 4->2, 5->3}
or
map{1->1, 2->4, 3->9, 4->2, 5->3}
```

efficiency of
comprehen-
sion
implementa-
tions

IMPLEMENTATION NOTE.

The fact that we define the meaning of generators inside comprehensions by replacing them with other constructs is not to suggest that this is a good implementation strategy. Even though it simplifies the language implementation, it is most likely very inefficient, introducing a lot of overhead in the form of closure creation, function application, and many intermediate collection objects that get added up to compute the final result.

6.11 Type assertion

A type assertion expression is a way to attach type information to an arbitrary expression inside the program code. Other than type assertions, the only expressions that the user explicitly provides type information for are variables, viz. when they are declared. The types of all other expressions are inferred

attaching type
information to
any
expression

from those. Depending on the type system, this inference may be more or less precise, and it may or may not be sufficient to guarantee a safe execution of the program code. In such cases, it might be useful to explicitly add a type assertion into the code, either to detect that it will not hold at compile time, or to be able to check it at runtime. The syntax looks as follows:

$$\text{TypeAssertionExpr} \rightarrow '(\text{ Expression } ':: \text{ Type })'$$

At compile time, if types are checked and inferred, there are three possible consequences of this construction depending on the type that the inference yielded for the embedded expression.

compile-time
type checking

1. The type can be shown to be a subtype of the asserted type. In this case, the assertion is always true.
2. The type can be shown to not intersect with the asserted type. In this case, objects computed by the expression will always fail to be of the asserted type, and thus this expression will always result in an error, which can be reported at compile time.
3. If neither can be shown, the expression may or may not produce objects of the asserted type. This results in three subcases:
 - (a) In a conservative system, an error will be signaled at compile time.
 - (b) The translator inserts a check that tests for compliance with the asserted type and causes an error if an object computed by the expression fails this check.
 - (c) The translator does nothing based on the assumption that the assertion implies that the expression will always produce the proper values.

runtime type
checks

Chapter 7

Statements

The execution of an action (as well as actor initialization) happens as the execution of a (possibly empty) sequence of statements. The only observable effect of a statement is a change of the variable assignments in its environment.¹ Consequently, the meaning of a statement is defined by how the variables in its scope change due to its execution. CAL provides the following kinds of statements:

statements
have “side-
effects”

Statement \rightarrow AssignmentStmt
| CallStmt
| BlockStmt
| IfStmt
| WhileStmt
| ForeachStmt
| ChooseStmt

7.1 Assignment

Assigning a new value to a variable is the fundamental form of changing the state of an actor. The syntax is as follows:

AssignmentStmt \rightarrow ID [Index | FieldRef] ':= ' Expression ';'
Index \rightarrow '[' [Expressions] ']'
FieldRef \rightarrow '.' ID

¹Other effects of statements, and indeed expressions, may include input/output operations, but we will disregard them in this context as they are not directly observable from within the language.

An assignment without an index or a field reference is a simple assignment, while one with a field reference is a field assignment, and one with an index is called an indexed assignment. Field assignments and indexed assignments are also collectively referred to as mutations (cf. sections 4.4 and 4.5 for more information on structured objects and mutability).

7.1.1 Simple assignment

In a simple assignment, the left-hand side is a variable name. A variable by that name must be visible in this scope, and it must be assignable.

The expression on the right-hand side must evaluate to an object of a value compatible with the variable (i.e. its type must be assignable to the declared type of the variable, if any—see section 4.1). The effect of the assignment is of course that the variable value is changed to the value of the expression. The original value is thereby overwritten.

assignability
→ 4.1

7.1.2 Field assignment

If a variable is of a type that has fields (see section 4.4), and if it is mutable (see section 4.5), assignments may also selectively assign to one of the fields rather than only to the variable itself. The syntax is as follows:

assigning to
field locations

FieldAssignmentStmt \rightarrow ID '.' ID ':=' Expression ','

Here, the first identifier is the variable name, while the second is the field name, which must be valid for the variable and the object that the variable refers to.

7.1.3 Assignment with indices

If a variable is of a type that is indexed, and if it is mutable, assignments may also selectively assign to one of its indexed locations, rather than only to the variable itself. The syntax is as follows:

assigning to
indexed
locations

IndexedAssignmentStmt \rightarrow ID '[' Expressions ']' ':=' Expression ','

In CAL, an indexed location inside an object is specified by a sequence of objects called indices, which are written after the identifier representing the variable, and which enclosed in square brackets.

indices

7.1.4 Assigning to and from mutable variables

In order to be able to reason about the actor state, and to facilitate program transformations, CAL is designed to avoid aliasing of stateful structures. In

no aliasing

other words, if a structure can be mutated, no two variables may point to it at the same time.

Therefore, when assigning a data structure to a mutable variable, that data structure must be cloned.² Of course, this cloning operation can occur on demand, or lazy, whenever the data structure, or a part of it, is mutated. Whichever implementation is chosen, mutations via a mutable variable must never have an effect on the value of other variables.

cloning

lazy cloning

Similarly, when assigning from a mutable variable, the structure assigned is conceptually copied, so that subsequent mutations of it are not visible via the new variable.

²In fact, it must be deeply cloned, up to the point where mutations can occur.

IMPLEMENTATION NOTE.

This may pose difficult implementation issues. Consider the following example:

```
mutable List[Integer] a = ...;
List[Integer] b := f(a[7, 1111]);
```

Let us assume that the indexer with two arguments on Lists computes the sublist from the first index to the second, inclusive, i.e. `a[7, 1111]` computes a list of length 1105 elements. A naive implementation would do just that, i.e. actually create the sublist. However, if the sublist is only an intermediate value in the computation of `f`, this would be very wasteful, e.g. in this case:

```
function f(List[Integer] v)-->List[Integer] :
  [#v]
end
```

Here, `f` returns a list of length 1 whose only element is the length of its parameter list.

Alternatively, a sublist could be represented by a special sublist-object that is backed by the original list, thus avoiding the explicit construction of the new structure. However, consider an `f` defined like this:

```
function f(List[Integer] v)-->List[Integer] :
  v
end
```

Now, `f` returns its parameter directly, with the consequence that whenever the original list, the one that is the value of the mutable variable `a`, is changed, so will the value of `b`, because its implementation is backed by the original list.

Obviously, an implementation that tries to achieve good performance therefore needs to do some bookkeeping of which parts of a data structure could be mutated, and when these get assigned to some variable, either clone them immediately, or mark them for cloning in case they should ever be mutated. In either case, the behavior must be as if the mutable data structure was cloned right away.

7.2 Procedure call

The only predefined operation on procedures (cf. section 6.9.2) is calling them, i.e. invoking them with a number of arguments. Calling a procedure is written as follows:

CallStmt \rightarrow Expression '(' Expressions ');'

The first expression must evaluate to a procedure, the other expressions must be of the appropriate argument types. The result of this statement is the execution of the procedure, with its formal parameters bound positionwise to the corresponding arguments.

7.3 Statement blocks (begin ... end)

Statement blocks are essentially syntactic sugar for a special case of the call statement, used to introduce a local scope and local variables. Their syntax looks like this:³

local scope

BlockStmt \rightarrow **begin** [**var** VarDecls **do**] { Statement } **end**

The form
`begin var decls do stmts end`
 is equivalent to the following procedure call:
`proc () var decls do stmts end () ;`

special case of
 procedure call

7.4 If-Statement

The if-statement is the most simple control-flow construct:

IfStmt \rightarrow **if** Expression **then** { Statement } [**else** { Statement }] **end**

As is to be expected, the statements following the **then** are executed only if the expression evaluates to **true**, otherwise the statements following the **else** are executed, if present. The expression must be of type `Boolean`.

7.5 While-Statement

Iteration constructs are used to repeatedly execute a sequence of statements. A **while**-construct repeats execution of the statements as long as a condition specified by a `Boolean` expression is true.

³Note that this is the one exception from the general rule that each keyword construct can end with an end marker that consists of the string `end` and the opening keyword. The keyword `endbegin` would have looked too awful.

WhileStmt \rightarrow **while** Expression [**var** VarDecls] **do** [Statements] (**end**|**endwhile**)

It is an error for the while-statement to not terminate.

7.6 Foreach-Statement

The **foreach**-construct allows to iterate over collections, successively binding variables to the elements of the collections and executing a sequence of statements for each such binding.

iteration over
collections

ForeachStmt \rightarrow ForeachGenerator { ',' ForeachGenerator }
 [**var** VarDecls] **do** [Statements] (**end**|**endforeach**)

ForeachGenerator \rightarrow **foreach** [Type] ID { ',' ID } **in** Expression
 [',' Expressions]

The basic structure and execution mechanics of the foreach-statement is not unlike that of the comprehensions with generators discussed in section 6.10.2. However, where in the case of comprehensions a collection was constructed piecewise through a number of steps specified by the generators, a foreach-statement executes a sequence of statements for each complete binding of its generator variables.

relation to
comprehen-
sions \rightarrow
6.10.2

Example 13. The following code fragment

```
s := 0;
foreach a in {1, 2}, b in {1, 2}:
  s := s + a*b;
end
```

results in *s* containing the number 9.

Foreach-statement semantics. In order to precisely describe the execution of a foreach-statement, we need to introduce a few symbols first. We call the original foreach-statement *S*. The (optional) declarations we write as *D*, and the body of the foreach-statement as *B*. Now *S* has $N \geq 1$ generators. The *i*-th generator has $M_i \geq 1$ variables, with names $v_{i,1}, \dots, v_{i,M_i}$. The collection expression of the *i*-th generator will be called *C_i*. Following the *i*-th generator are $K_i \geq 0$ filter expressions, which we call $F_{i,1}, \dots, F_{i,K_i}$.

We can now define the meaning of the foreach-statement by replacing it with a number of previously described constructs, such as procedure closures, procedure call, and if-statement. The key to this replacement is a function `$iterate`, which we will describe below.

```
foreach
  v1,1, ..., v1,M1
in C1,
  F1,1, ..., F1,K1,
...
foreach
  vn,1, ..., vn,Mn
in Cn,
  Fn,1, ..., Fn,Kn
var D do B
end
syntactical
transforma-
tion of
foreach-
statement
```

$$\begin{aligned}
S &\equiv G(1) \\
G(i) &\equiv \begin{cases} \underline{\text{begin var } D \text{ do } B \text{ end}} & i > N \\ \underline{\text{begin var } a_i \equiv C_i \text{ do } G_V(i, 1) \text{ end}} & \text{otherwise} \end{cases} \\
G_V(i, j) &\equiv \begin{cases} G_F(i, 1) & j > M_i \\ \underline{\$iterate}(a_i, \underline{\text{proc } (v_{i,j}) \text{ do } G_V(i, j + 1) \text{ end}}) & \text{otherwise} \end{cases} \\
G_F(i, k) &\equiv \begin{cases} G(i + 1) & k > K_i \\ \underline{\text{if } F_{i,k} \text{ then } G_F(i, k + 1) \text{ end}} & \text{otherwise} \end{cases}
\end{aligned}$$

The a_i are distinct and fresh variable symbols, i.e. they are mutually different, and they also differ from any variable symbol occurring inside the foreach-statement.

The `$iterate` procedure takes two arguments, a collection and a unary procedure. It iterates over the elements of its first argument, and calls the procedure on each element. `$iterate` is defined in appendix C.2.1.

See also page 46 for an explanation of the iterator name choice, and page 46 for a note concerning the efficiency of generator implementations via syntactic substitution.

`$iterate:`
`Collection[T]`
`[T-->]`
`Def. → C.2.1`

7.7 Choose-Statement

The choose-statement permits the explicit specification of (potentially non-deterministic) choice among a set of alternatives.

```

ChooseStmt → ChooseGenerator { ',' ChooseGenerator }
           [var VarDecls] do [Statements]
           [else[[var LocalVarDecls] do] [Statements] ] (end|endchoose)
ChooseGenerator → choose [Type] ID [',' IDs] in Expression [',' Expressions]

```

The binding of values to variables happens very much like in the case of comprehension generators or the foreach-statement—see below for a precise definition of the semantics.

In contrast to the foreach-statement, the choose-statement executes its body at most once, viz. for the first combination of values that satisfy all the filters (which are part of the generators just as in the case of the foreach-statement or collection comprehensions). If no such combination can be found, and if an else-branch is present, the statements following the **else**-keyword are executed instead.

Choose-statement semantics. In order to precisely describe the execution of a choose-statement, we need to introduce a few symbols first. We call the

original choose-statement S . The (optional) declarations we write as D , and the body of the foreach-statement as B , the else branch is E . Now S has $N \geq 1$ generators. The i -th generator has $M_i \geq 1$ variables, with names $v_{i,1}, \dots, v_{i,M_i}$. The collection expression of the i -th generator will be called C_i . Following the i -th generator are $K_i \geq 0$ filter expressions, which we call $F_{i,1}, \dots, F_{i,K_i}$.

We can now define the meaning of the choose-statement by replacing it with a number of previously described constructs, such as procedure closures, procedure call, and if-statement. The key to this replacement is a function $\$try$, which we will describe below.

$$\begin{aligned}
 S &\equiv \underline{\text{begin var } a_* := \text{false do } G(1) \text{ if not } a_* \text{ then } E \text{ end end}} \\
 G(i) &\equiv \begin{cases} \underline{\text{begin var } D \text{ do } a_* := \text{true; } B \text{ end}} & i > N \\ \underline{\text{begin var } a_i \equiv C_i \text{ do } G_V(i, 1) \text{ end}} & \text{otherwise} \end{cases} \\
 G_V(i, j) &\equiv \begin{cases} G_F(i, 1) & j > M_i \\ \underline{\$try(a_i, \text{proc } (v_{i,j}) \text{ do } G_V(i, j+1) \text{ end, lambda() : } a_* \text{ end})} & \text{otherwise} \end{cases} \\
 G_F(i, k) &\equiv \begin{cases} G(i+1) & k > K_i \\ \underline{\text{if } F_{i,k} \text{ then } G_F(i, k+1) \text{ end}} & \text{otherwise} \end{cases}
 \end{aligned}$$

```

choose
  v1,1, ..., v1,M1
in C1,
  F1,1, ..., F1,K1,
  ...
choose
  vn,1, ..., vn,Mn
in Cn,
  Fn,1, ..., Fn,Kn
var D do B
else E end
syntactical
transformation of
choose-statement

```

The a_i and a_* are distinct and fresh variable symbols, i.e. they are mutually different, and they also differ from any variable symbol occurring inside the choose-statement.

The $\$try$ procedure takes three arguments, a collection, a unary procedure, and a nullary function. It iterates over the elements of its first argument, and for each element it does the following: It applies the function (to no arguments, as it is nullary). If the value of that application is **false**, it proceeds to call the procedure (its second argument) on the element, otherwise it simply returns (disregarding any subsequent elements of the collection).

```

$try:
  Collection[T]
  [T-->]
  [-->Boolean]
Def. → C.2.1

```

See also page 46 for an explanation of the iterator name choice, and page 46 for a note concerning the efficiency of generator implementations via syntactic substitution. $\$try$ is defined in appendix C.2.1.

RATIONALE.

Having a construct that explicitly allows the use of non-determinism (as opposed to expressions, such as map comprehensions, that may evaluate nondeterministically, but where this behavior is most likely unwanted and potentially erroneous) allows actor authors to express internal choice points which are explicitly not under the control of the actor context (i.e. the model of computation).

The responsibility for many other choices, such as which action to fire, can be assumed by the model of computation. Using a choose-statement inside an action provides an explicit signal that this action may be nondeterministic, and allows the model of computation to either reject it, or deal with it accordingly.

[backtracking](#)

Example 14. The choose-statement can be used to implement simple backtracking algorithms, as in the following example:

```
s := null;
choose
  a in [1, 2],
  choose
    b in if a = 1 then {} else {3, 4} end
  do
    s := [a, b];
  else
    s := [];
  end
end
```

Since for $a = 1$ the collection for b is the empty set, the body will never be executed for this value of a , because no choice could be made for b . However, if $a = 2$, two choices can be made for b , and it is unspecified which one is made. After executing this statement, s will be either $[2, 3]$ or $[2, 4]$.

It will never be $[\]$, because the else branch is only executed if no choice can be made.

Chapter 8

Actions

An action in CAL represents a (often large or even infinite) number of transitions of the actor transition system described in section 10.3. A CAL actor description can contain any number of actions, including none. The definition of an action includes the following information:

action \equiv
family of
transitions

- its input tokens,
- its output tokens,
- the state change of the actor,
- additional firing conditions,
- the time delay.

In any given state, an actor may take any number of transitions (including zero), and these transitions may be represented by any number of actions in the actor description. The choice between them is ultimately made by the context of an actor (though the actor can constrain the possible choices, see chapter 9)—see section 8.4 for details.

choice
between
actions

The syntax of an action definition is as follows:

```
Action  $\rightarrow$  [ActionTag ':' ] action ActionHead [do Statements] (end|endaction)  
ActionTag  $\rightarrow$  ID { '.' ID }  
ActionHead  $\rightarrow$  InputPatterns '==>' OutputExpressions  
[guard Expressions ] [var VarDecls ] [delay Expression]
```

Actions are optionally preceded by action tags which come in the form of qualified identifiers (i.e. sequences of identifiers separated by dots), see also section 9.1. These tags need not be unique, i.e. the same tag may be used for more than one action. Action tags are used to refer to actions, or sets of actions, in action schedules and action priority orders—see chapter 9 for details.

action tags

The head of an action contains a description of the kind of inputs this action applies to, as well as the output it produces. The body of the action is a sequence of statements that can change the state, or compute values for local variables that can be used inside the output expressions.

Input patterns and output expressions are associated with ports either by position or by name. These two kinds of association cannot be mixed. So if the actor's port signature is

```
Input1, Input2 ==> ...
```

an input pattern may look like this:

```
[a], [b, c]
```

(binding `a` to the first token coming in on `Input1`, and binding `b` and `c` to the first two tokens from `Input2`). It may also look like this:

```
Input2: [c]
```

but never like this:

```
[d] Input2:[e]
```

This mechanism is the same for input patterns and output expressions.

The following sections elaborate on the structure of the input patterns and output expressions describing the input and output behavior of an action, as well as the way the action is selected from the set of all actions of an actor.

In discussing the meaning of actions and their parts it is important to keep in mind that the interpretation of actions is left to the model of computation, and is not a property of the actor itself. It is therefore best to think of an action as a declarative description of how input tokens, output tokens, and state transitions are related to each other. See also section 8.4.

actions
declarative

8.1 Input patterns, and variable declarations

Input patterns, together with variable declarations and guards, perform two main functions: (1) They define the input tokens required by an action to fire, i.e. they give the basic conditions for the action to be firable which may depend on the value and number of input tokens and on the actor state, and (2) they declare a number of variables which can be used in the remainder of the action to refer to the input tokens themselves. This is their syntax:

input pattern:
activation
condition &
variable
declaration

```
InputPattern → [ID ':' ] '[' IDs ']' [RepeatClause] [ChannelSelector]
```

```
ChannelSelector → at Expression
```

```
| at* Expression
```

```
| [at*] any
```

```
| [at*] all
```

```
RepeatClause → repeat Expression
```

Input patterns differ depending on whether they are written for a single port or a multiport—the former represents one channel of incoming tokens,

multiport
pattern
—channel
selector

	no repeat-clause	with repeat-clause
single channel	T	List[T]
multichannel	Map[<i>CID</i> , T]	Map[<i>CID</i> , List[T]]

Table 8.1: Token variable types depending on input pattern kind and presence of repeat-clause. (T is the token type of the corresponding input port.)

while the latter represents an arbitrary number (including zero) of channels. A multiport input pattern has to have a channel selector associated with it, which is a construction that serves to specify the channels that a pattern is applied to (i.e. that the tokens bound by the pattern are read from).

The static type of the variables declared in an input pattern depends on the token type declared on the input port, but also on the kind of pattern. Fig. 8.1 shows this dependency. It distinguishes between patterns according to whether they contain a repeat-clause, and whether they take tokens from one channel (single-channel patterns), or any number of channels (multichannel patterns). Channels are specified using channel identifiers, which are objects identifying a specific channel inside a multiport. Usually, channel identifiers are represented by some other type in some platform-dependent manner. For instance, non-negative integer numbers can be used to identify channels. Fig. 8.1 refers to the type of channel identifiers as *CID*.

8.1.1 Single-port input patterns

A single-port input pattern binds a number of tokens from the input channel associated with the specified port to variables specified as part of the pattern. Because there is exactly one input channel associated with a single port, these patterns are always single-channel patterns.

single-port
 \Rightarrow
 single-channel

SinglePortInputPattern \rightarrow [ID ':'] '[IDs]' [RepeatExpression]

A pattern without a repeat-expression is just a number of variable names inside square brackets. The pattern binds each of the variable names to one token, reading as many tokens as there are variable names. The number of variable names is also referred to as the pattern length. The static type of the variables is the same as the token type of the corresponding port (Fig. 8.1).

Example 15 (Single-port input pattern). Assume the sequence of tokens on the input channel is the natural numbers starting at 1, i.e.

1, 2, 3, 4, 5, ...

The input pattern [a, b, c] results in the following bindings:

a = 1, b = 2, c = 3

If the pattern contains a repeat-clause, that expression must evaluate to a non-negative integer, say *N*. If the pattern length is *L*, the number of tokens

read by this input pattern and bound to the L pattern variables is NL . Since in general there may be more tokens than variables (N times more, exactly), variables are bound to lists of tokens, each list being of length N . In the pattern, the list bound to the k -th variable contains the tokens numbered $k, L+k, 2L+k, \dots, (N-1)L+k$. The static type of these variables is `List[T]`, where `T` is the token type of the port (Fig. 8.1).

Example 16 (Single-port input pattern with repeat-clause). Assume again the natural numbers as input sequence. If the input pattern is

```
[a, b, c] repeat 2
```

it will produce the following bindings:

```
a = [1, 4], b = [2, 5], c = [3, 6]
```

8.1.2 Multiport input patterns

Multiport input patterns contain an expression that functions as a channel selector, i.e. it specifies which channels to read tokens from. There are two kinds of channel selectors—those that identify precisely one channel, and those that identify any number of channels.

channel selector: one or many

Single-channel channel selectors. The `at`-construct is used to select one channel from a multiport. The expression following the `at`-keyword must evaluate to a valid channel identifier for the given port. The reading of tokens and binding to the pattern variables happens in the same way as described for single-port input patterns above.

Example 17 (at channel selector). Assume four input channels `'a'`, `'b'`, `'c'`, `'d'` with the following input tokens:

```
'a': 11, 12, 13, 14, 15, 16, 17, 18, 19
```

```
'b': 21, 22, 23, 24, 25, 26, 27, 28, 29
```

```
'c': 31, 32, 33, 34, 35, 36, 37, 38, 39
```

```
'd': 41, 42, 43, 44, 45, 46, 47, 48, 49
```

Then the input pattern

```
[a, b, c] repeat 2 at 'c'
```

would yield the bindings

```
a = [31, 34], b = [32, 35], c = [33, 36]
```

Multichannel channel selectors. CAL provides three constructs for selecting tokens from more than one channel of a multiport at the same time. They are distinguished by the keyword that introduces them: `at*`, `any`, `all`. They differ in the way they determine the set of channels to read from as follows:

different ways to determine channel set

1. Following the `at*`-keyword is an expression that computes a (possibly empty) collection of channel identifiers.

2. The **any**-keyword results in those channels being read from that have a sufficient number of tokens available to be bound to the variables of the pattern—which could be none at all, in which case the set of channels would be empty.
3. The **all**-keyword selects all channels of the multiport.

All multichannel input patterns have in common that they are applied homogeneously, i.e. the same number of tokens is read from all selected channels (and no tokens are read from any of the other channels). The variables are bound to maps—the keys of these maps are channel identifiers, while the values are, depending on the absence or presence of a repeat-clause, either tokens or lists of tokens. Fig. 8.1 gives the static types of the pattern variables for multichannel patterns.

Example 18 (Multichannel input patterns). Assume four input channels 'a', 'b', 'c', 'd' with the following input tokens:

```
'a': 11, 12, 13, 14, 15, 16, 17, 18, 19
'b': 21, 22, 23, 24, 25
'c': 31, 32, 33, 34, 35, 36, 37, 38, 39
'd': 41, 42, 43
```

The input pattern

```
[a, b, c] at* {'a', 'c'}
```

yields the bindings

```
a = map{'a'->11, 'c'->31}
b = map{'a'->12, 'c'->32}
c = map{'a'->13, 'c'->33}
```

The input pattern

```
[a, b] all
```

results in the bindings

```
a = map{'a'->11, 'b'->21, 'c'->31, 'd'->41}
b = map{'a'->12, 'b'->22, 'c'->32, 'd'->42}
```

In this case, the keyword **any** instead of **all** would have produced the same result.

The input pattern

```
[a, b] repeat 2 any
```

results in the bindings

```
a = map{'a'->[11, 13], 'b'->[21, 23], 'c'->[31, 33]}
b = map{'a'->[12, 14], 'b'->[22, 24], 'c'->[32, 34]}
```

In this case, **all** instead of **any** would not have worked, because the channel 'd' does not provide the required number of tokens.

8.1.3 Scoping of action variables

The scope of the variables inside the input patterns, as well as the explicitly declared variables in the var-clause of an action is the entire action—as a con-

sequence, these variables can depend on each other. The general scoping rules from section 5.2 need to be adapted in order to properly handle this situation.

In particular, input pattern variables do not have any initialization expression that would make them depend explicitly on any other variable. However, their values clearly depend on the expressions in the repeat-clause and the channel selector (if present). For this reason, for any input pattern variable v , we define the set of free variables of its initialization expression F_v to be the union of the free variables of the corresponding expressions in the repeat-clause and the channel selector.

The permissible dependencies then follow from the rules in section 5.2.

Example 19 (Action variable scope). The following action skeleton contains dependencies between input pattern variables and explicitly declared variables

```
[n] at c, [k], [a] repeat m * n ==> ...
var
  m = k * k, c = f(m)
do ... end
```

These declarations are well-formed, because the variables can be evaluated in the order k, m, c, n, a .

By contrast, the following action heads create **circular dependencies** and are therefore **errors**:

```
[n] at f(a), [a] repeat n ==> ... do ... end

[a] repeat a[0] + 1 ==> ... do ... end

[a] repeat n ==> ...
var
  n = f(b), b = sum(a)
do ... end

[a] at c, [b] repeat a ==> ...
var
  c = g(b)
do ... end
```

8.2 Output expressions

Output expressions are conceptually the dual notion to input patterns—they are syntactically similar, but rather than containing a list of variable names which get bound to input tokens they contain a list of expressions that computes the output tokens, the so-called token expressions.

pattern
variables \leftrightarrow
var-variables
 F_v for pattern
variables
 \rightarrow 5.2, Def. 1

output
expressions
dual of input
patterns

	no repeat-clause	with repeat-clause
single channel	T	Seq[T]
multichannel	Map[CID, T]	Map[CID, Seq[T]]

Table 8.2: Token expression types depending on output expression kind and presence of a repeat-clause. (T is the token type of the corresponding output port.)

OutputExpression \rightarrow [ID ':'] '[' Expressions ']' [RepeatClause] [ChannelSelector]

ChannelSelector \rightarrow **at** Expression

| **at*** Expression
 | [**at***] **any**
 | [**at***] **all**

RepeatClause \rightarrow **repeat** Expression

The repeat-clause and channel selector work not unlike in the case of input patterns, but with one crucial difference. For input patterns, these constructs control the construction of a data structure that was assembled from input tokens and then bound the pattern variables. In the case out output expressions, the values computed by the token expressions are themselves these data structures, and they are indexed into according to the repeat-clause and the channel selector, if these are present.

The table in Fig. 8.2 shows the kind of data structure a token expression has to evaluate to depending on the presence or absence of a repeat-clause, and depending on whether the output expression is a single-channel or a multichannel output expression.

Single-channel output expressions. In single-channel output expressions without repeat-clause, the token expressions represent the output tokens directly, and the number of output tokens produced is equal to the number of token expressions. If a single-channel output expression does have a repeat-clause, the token expressions must evaluate to sequences of tokens, and the number of tokens produced is the product of the number of token expressions and the value of the repeat-expression. In addition, the value of the repeat-expression is the minimum number of tokens each of the sequences must contain.

Example 20 (Single-channel output expressions.). The output expression

```
... ==> [1, 2, 3]
```

produces the output tokens 1, 2, 3.

The output expression

```
... ==> [[1, 2, 3], [4, 5], [6, 7, 8]] repeat 2
```

produces the output tokens 1, 4, 6, 2, 5, 7.

deconstructing
output
expression
values

repeat-clause:
lists must be
long enough!

The output expression
`... ==> [[1, 2, 3], [4, 5]] repeat 1 at 'b'`
 produces the output tokens 1, 4 on channel 'b'.

RATIONALE.

One of the key reasons for interpreting the repeat-construct in output expressions as described above is to maintain symmetry with the way input patterns are handled. For instance, consider the following action:

```
action [a, b, c] repeat k ==> [a, b, c] repeat k end
```

This action copies its input tokens identically to its output, $3k$ tokens at a time. It is, therefore, equivalent to the action

```
action [a] repeat 3 * k ==> [a] repeat 3 * k end
```

The same rationale applies to multichannel output expressions, see below. Here, too, the design strives for a symmetry between the input patterns and output expressions, as far as their fundamental directedness allows (the former declares variables, while the latter uses them to compute values).

Multichannel output expressions. Similarly, for multichannel output expressions without a repeat expression, the token expressions must evaluate to maps from channel identifiers to tokens. If a repeat-clause is present, the maps must map channel identifiers to sequences of tokens, the number of elements on those sequences being not smaller than the value of the repeat expression. The channel selector places some constraints on the domains of the maps as follows:

- For an **at*** channel selector, the intersection of the domains of the maps must contain the collection of channel identifiers the selector expression evaluates to.
- For an **all** channel selector, the intersection of all domains must contain the set of all channel identifiers currently valid for the corresponding port.
- For a **any** channel selector, no constraints apply to the domains. The channels that output will be produced on are those in the intersection of all domains, which may be empty.

These rules are designed to make it possible to produce output homogeneously on all selected output channels—in other words, all output channels of a multiport output the same number of tokens if they output any tokens at all.

Example 21 (Multichannel output expressions.). The output expression

```
... ==> [map {'a'->1, 'b'->2},
         map {'a'->3, 'b'->4, 'c'->5}] at* {'a', 'b'}
```

channel
selector
constrains
token map
domains

homogeneous
output
production

produces the following output on the channels 'a' and 'b':

```
a: 1, 3
b: 2, 4
```

Nothing is output on channel 'c'.

The output expression

```
... ==> [map {'a'->[1, 2], 'b'->[3, 4]},
         map {'a'->[5, 6], 'b'->[7, 8]}]
         repeat 2 at* {'a', 'b'}
```

results in the following output:

```
a: 1, 5, 2, 6
b: 3, 7, 4, 8
```

8.3 Delays

The expression following the **delay**-keyword specifies the delay of the action. The actor language itself does not attach any meaning to this information, other than that it is part of the label of a state transition of an actor. See section 10.2 for the precise meaning of time and delay, and section 3.2 for some more information on time inside the CAL language.

uninterpreted
delays

If the actor contains a time-clause, then the static type of each delay expression must be a subtype of the type specified in that time-clause.

8.4 On action selection: guards and other activation conditions

At any given point during the execution of an actor, an action may potentially fire on some input data. Whether it is activated, i.e. whether in a given situation it actually can fire, depends on whether its activation conditions have all been met. The minimal conditions are as follows:

activation
conditions

1. According to the action schedule (see section 9.2) this action may fire next, according to Def. 4 on page 70.
2. No higher-priority action is activated (see section 9.3).
3. There are sufficient input tokens available to bind the input pattern variables to appropriate values.
4. Given such a binding, all guard expressions (which must be Boolean expressions) evaluate to **true**.
5. There is sufficient room for the output tokens produced by this firing.

Depending on the context in which the actor is embedded, additional conditions for activating actions may apply. A model of computation may reject an actor as non-well-formed if it cannot statically ascertain that it will meet the relevant activation conditions whenever it needs to.

additional
conditions by
model of
computation

IMPLEMENTATION NOTE.

The last activation condition may not be easily satisfiable in the general case, as it may be unknown exactly how many output tokens a firing will produce until it has executed and the output tokens have been computed. The important point is that firing an actor has to appear to the outside as an atomic step—the physical time it takes is not directly relevant, and the execution may even stall due to a lack of space for the output tokens, but it is imperative that the firing looks as if it occurred in one atomic transition.

8.5 Initialization actions

Initialization actions are executed at the beginning of an actor’s life cycle. They are very similar to regular actions, with two important differences:

special action:
no input, no
invariants

1. Since the assumption is that at the beginning of an actor execution no input is available, initialization actions have no input patterns. They may produce output, however.
2. With the exception of initialization expressions in variable declarations, an initialization action contains the first code to be executed inside the actor. Any state invariants in the actor may not hold, and instead have to be established by the initialization action.

The syntax of initialization actions is as follows:

```
InitializationAction → [ActionTag ':']
    initialize InitializerHead [do Statements] (end|endinitialize)
InitializerHead → '==>' OutputExpressions
    [guard Expressions ] [var VarDecls ] [delay Expression]
```

The activation conditions for actions apply also to initialization actions—of course, since there is no input, the conditions concerning input tokens become vacuously true.

activation
conditions

If an actor has more than one initialization action, and if more than one is activated at the beginning of an actor execution, one of them is chosen arbitrarily. The actor context can of course make than choice.

only one init
action gets
fired

Chapter 9

Action-level control structures

In CAL, an action expresses a relation between the state of an actor and input tokens, and the successor state of the actor and output tokens. In general, CAL actors may contain any number of actions, and in a given situation, any subset of those may be ready to be executed. For example, both actions of the following actor may be able to execute, if there is a token available on either input port:

Example 22 (Nondeterministic Merge).

```
actor NDMerge () A, B ==> C :  
  
  action A: [x] ==> [x] end  
  
  action B: [x] ==> [x] end  
end
```

ambiguous
action choice

It is important to emphasize that the policy used to choose between the two actions above is not part of the actor specification. In practice, the context of the actor will make some choice, based on whatever policy it implements. This flexibility may be desired, but sometimes the actor writer may want to have more control over the choice of the action—e.g., if the Merge actor is supposed to alternate between reading its input ports, one might use actor state to realize this behavior:

control over
action
selection

Example 23 (Basic FairMerge).

```
actor FairMerge () A, B ==> C :  
  
  s := 1;  
  
  action A: [x] ==> [x]  
  guard s = 1  
  do  
    s := 2;
```

using state to
control action
selection

```

end

action B: [x] ==> [x]
guard s = 2
do
  s := 1;
end
end
end

```

This way of specifying action choice has two key drawbacks. First, it is very cumbersome to write and maintain, and it does not scale very well even for modest numbers of actions and states. Furthermore, this way of specifying action choice essentially obfuscates the "real" logic behind guards, state variables, and assignments, so that it becomes harder to extract the intent from the actor description, both for tools and for human readers.

schedules
easier to read
and to analyze

These are the key motivations for using action schedules, i.e. structured descriptions of possible orders in which actions may fire. Before we can discuss action schedules in section 9.2, we need to take a closer look at how actions are referred to inside of them.

9.1 Action tags

Actions are optionally prefixed with action tags (see chapter 8), which are qualified identifiers:

```

ActionTag → QualID
QualID → ID { '.' ID }

```

The same tag may be used for more than one action. In the following, we write the set of all actions tagged by a tag t as \bar{t} , and the tag of some action a as t_a . The empty tag is written as ϵ , and the set of all untagged actions is therefore $\bar{\epsilon}$.

multiple
actions per tag

Action tags are ordered by a prefix ordering: We say that $t \sqsubseteq t'$, i.e. t is a prefix of t' , if t' starts with all the identifiers in t in the same order, followed by any number of additional identifiers, including none. For instance, $a.b.c \sqsubseteq a.b.c.x$ and $a.b \sqsubseteq a.b$, but $a.b \not\sqsubseteq a.c$. We call t' an extension of t .

prefix order
among tags

When used inside action schedules and priority orderings, a tag denotes the set of actions which are labeled with tags that are extensions of it. For any tag t this set is called \hat{t} and is defined as follows:

tag: denoted
action set

$$\hat{t} =_{def} \{a \mid t \sqsubseteq t_a\}$$

9.2 Action schedules

Action schedules are structured descriptions of possible sequences in which the actions of an actor may be executed. These sequences are either specified by a finite state machine or a regular expression in the alphabet of action tags. In general, like any regular language, the set of possible sequences may be finite or infinite.

fsm & regexp
schedules

Irrespective of whether a finite state machine is used or a regular expression, an action schedule effectively describes a (regular) language \mathcal{L} in the alphabet of action tags. This language is used to constrain the legal sequences of action firings as follows.

schedule =
regular tag
language

Definition 4 (Legal action sequence). Given a tag language \mathcal{L} , assume a finite sequence of actions $(a_i)_{i=1..n}$, and a sequence $(b_j)_{j=1..m}$ with $m \leq n$ and a strict monotonic function $f : \{1..m\} \rightarrow \{1..n\}$ such that the following holds for all $j \in \{1..m\}$ and $i \in \{1..n\}$:

1. $b_j = a_{f(j)}$
2. $t_{b_j} \neq \epsilon$
3. $t_{a_i} = \epsilon, \quad \forall i \notin f^{-1}[\{1..m\}]$

(b_j) tagged
subsequence
of (a_i)

In other words, the (b_j) are the subsequence in the (a_i) with non-empty tags. If (b_j) is empty, i.e. $m = 0$, then (a_i) is a *legal action sequence*.

If (b_j) is not empty, i.e. $m \geq 1$, then (a_i) is a *legal action sequence* if and only if there exists a sequence of tags $(t_j)_{j=1..m}$ such that the following holds:

1. for all $j \in \{1..m\}$, $b_j \in \hat{t}_j$
2. there exists a $w \in \mathcal{L}$ such that $(t_j) \sqsubseteq w$.

A consequence of this definition is that untagged actions may occur at any point in the schedule—conversely, schedules do not constrain untagged actions in any way.

no constraints
on untagged
actions

We will now describe the two ways of defining a tag language: finite state machines and regular expressions.

9.2.1 Finite state machine schedules

A finite state machine schedule defines a number of transitions between states (and an initial state) that are each labeled with one or more action tags.

```
ScheduleFSM → schedule [fsm] ID ':'
  { StateTransition ';' }
  (end|endschedule)
```

```
StateTransition → ID '(' ActionTags ')' '- ->' ID
  { '|' '(' ActionTags ')' '- ->' ID }
```

```
ActionTags → ActionTag
           | ActionTag ',' ActionTags
```

The state before the colon is the initial state, and all states are accepting (or final states). The tag language is the set of all sequences of tags that label transitions leading from the initial state to any other state of the finite state machine.

Several transitions starting from the same state may be written as separated by the ',' character.

The following illustrates the use of a finite state machine action schedule to express the `FairMerge` actor somewhat more concisely.

Example 24 (FairMerge, with FSM schedule).

```
actor FairMerge1 () A, B ==> C :

  InA: action A: [x] ==> [x] end

  InB: action B: [x] ==> [x] end

  schedule fsm WaitA :
    WaitA (InA) --> WaitB;
    WaitB (InB) --> WaitA;
  end
end
```

9.2.2 Regular expression schedules

In many cases, using an explicit finite state machine to define an action schedule is still rather verbose, and specifying the tag language through a regular expression is much more straightforward. The syntax is as follows:

```
ScheduleRegExp → schedule regexp RegExp (end|endschedule)
RegExp → ActionTag
       | '(' RegExp ')'
       | '[' RegExp ']'
       | RegExp '*'
       | RegExp RegExp
       | RegExp '|' RegExp
```

The simplest regular expression (`regexp`) is simply a tag, which denotes the language consisting of just this tag. Parentheses are used to explicitly group regexps. The square brackets enclose an optional regexp, describing the language that consists of all sequences described by the `regexp`, as well as the empty

sequence. The Kleene operator '*' means repetition any number of times (including zero), juxtaposition of two regexps denotes concatenation of sequences of the two languages, and the '|' operator describes alternative choice. The binding strength is the order given above, i.e. the regexp

$$A B^* \mid C^* D$$

is interpreted as

$$(A (B^*)) \mid ((C^*) D)$$

The FairMerge actor can be expressed with regexps as follows:

Example 25 (FairMerge, with regular expression schedule).

```
actor FairMerge2 () A, B ==> C :
  InA: action A: [x] ==> [x] end
  InB: action B: [x] ==> [x] end

  schedule regexp
    (InA InB)*
  end
end
```

operator
precedence

RATIONALE.

The reason for this conceptual redundancy is convenience—many schedules are best written as regular expressions. On the other hand, in many cases producing a regular expression can be cumbersome, and finite state machines are preferable. This is especially true for generated code, which will usually use the finite state machine formulation.

Of course, both syntaxes are of identical expressiveness, and can therefore be translated into one another.

Using regular expressions, control structures among actions become much more concise and easier to read and to specify. For example, assume we do not want to predetermine which port of the FairMerge actor is read from first. It is very easy to augment the FairMerge2 description to realize this:

Example 26 (Symmetric FairMerge).

```

actor FairMergeSymmetric () A, B ==> C :

  InA: action A: [x] ==> [x] end

  InB: action B: [x] ==> [x] end

  schedule regexp
    (InA InB)* | (InB InA)*
  end
end

```

9.3 Priorities

Priorities are very different from action schedules in that they genuinely add to the expressiveness of CAL—it would not be possible in general to reduce them to existing constructs, in the way schedules can in principle be reduced to a state variable and guards/assignments. Among other things, priorities allow actors to effectively test for the absence of tokens. As a consequence, actors can express non-prefix monotonic processes,¹ which is powerful, but at the same time can be dangerous, because it means the results computed by an actor may depend on the way it was scheduled with respect to the other actors in the system.

Priorities are defined as a partial order relation over action tags, which induces a partial order relation over the actions. An action can only fire if there is no other enabled action that is higher in this partial order. The order is specified as follows:

```

PriorityOrder → priority{ PriorityInequality ';' } end
PriorityInequality → ActionTag '>' ActionTag { '>' ActionTag }

```

The priority inequalities are specified over tags, i.e. they have the form $t_1 > t_2$. These inequalities induce a binary relation on the actions as follows:

$$a_1 > a_2 \iff \exists t_1, t_2 : t_1 > t_2 \wedge a_1 \in \hat{t}_1 \wedge a_2 \in \hat{t}_2 \\ \vee \exists a_3 : a_1 > a_3 \wedge a_3 > a_2$$

The priority inequalities are valid iff the induced relation $>$ on the actions is an irreflexive partial order, i.e. it is antisymmetric and transitive. Transitivity follows from the definition, but antisymmetry and irreflexivity do not. In fact, they do not even follow if the relation on the tags is a partial order. Consider

¹See [10] for details on prefix monotonicity and its implications for the description of dataflow systems.

priorities
increase ex-
pressiveness

allow test for
absence of
tokens

induced
partial order
on actions

validity of
priority
system

the following example:

```
A.B > X > A
```

This is obviously a proper order on the tags. However, if we have two actions labeled X and A.B, then the induced relation is not antisymmetric, because according to $A.B > X$, the action A.B has higher priority than X, while the inequality $X > A$ implies that any action whose label starts with X has a higher priority than any action whose label starts with A, including, of course, A.B. Therefore, this system of priority inequalities is inconsistent and thus invalid.

With priorities, we can express a Merge actor that prefers one input over the other like this:

Example 27 (BiasedMerge).

```
actor BiasedMerge () A, B ==> C :

  InA: action A: [x] ==> [x] end
  InB: action B: [x] ==> [x] end

  priority
    InA > InB;
  end
end
```

Perhaps more interestingly, we can express a merge actor that is fair, in the sense that it will consume equal amounts of tokens from both inputs as long as they are available, but will not halt due to lack of tokens on only one of its input ports. It is also nondeterministic, i.e. it does not specify the order in which it outputs the tokens.

Example 28 (FairMerge, with priorities).

```
actor FairMerge3 () A, B ==> C :

  Both: action [x], [y] ==> [x, y] end
  Both: action [x], [y] ==> [y, x] end
  One:  action A: [x] ==> [x] end
  One:  action B: [x] ==> [x] end

  priority
    Both > One;
  end
end
```

Part II

Semantics

Chapter 10

Actor model

This chapter presents the formal definition of our notion of actor. In general, actors may be highly complex and infinite computational structures, which can be non-deterministic, contain and manipulate internal state, and consume and produce units of data (tokens). The purpose of the CAL language (as that of any other programming language) can be seen as providing a finite representation of such structures, while exhibiting their inner structures and regularities in a way that allows tools to identify and check them, and to use them when composing actors or when generating implementations from actor descriptions.

language:
finite
representation
of infinite
structure

10.1 Preliminaries

In the following, we assume a universe of all token values \mathcal{U} that can be exchanged between actors. The communication along each connection between actors can be viewed as a sequential stream of tokens, and actors will remove tokens from this stream and add tokens to it. We define the set $\mathbb{S} =_{def} \mathcal{U}^*$ as the set of all finite sequences over \mathcal{U} . The set $\mathbb{S}_\infty =_{def} \mathbb{S} \cup \mathbb{S}^\mathbb{N}$ is the set of all finite and infinite sequences over \mathcal{U} . We write the empty sequence as λ . The length of a finite sequence s is denoted by $|s|$.

universe \mathcal{U}
finite
sequences \mathbb{S}
infinite
sequences \mathbb{S}_∞

The elements of \mathbb{S}_∞ (and consequently \mathbb{S}) are partially ordered by their prefix relation: $s \sqsubseteq r$ iff s is a prefix of r , i.e. r starts with the tokens of s in the order they occur in s . For example, $abc \sqsubseteq abcd$, and $ab \sqsubseteq ab$, but $abc \not\sqsubseteq cabd$. Note that for any $s \in \mathbb{S}_\infty$, $\lambda \sqsubseteq s$ and $s \sqsubseteq s$.

prefix order

Many actors have multiple input and output sequences, so most of the time we will work with tuples of sequences, i.e. with elements of \mathbb{S}^n or \mathbb{S}_∞^n for some n . The prefix order extends naturally to tuples as follows:

... extends to
 \mathbb{S}^n

$$(s_i)_{i=1..n} \sqsubseteq (r_i)_{i=1..n} \Leftrightarrow_{def} \forall i = 1..n : s_i \sqsubseteq r_i$$

Note that since \mathbb{S}_∞ and \mathbb{S}_∞^n under the prefix order both have a least element, and every chain in them has a least upper bound, they are complete partial

orders. This property is relevant when relating dataflow models with a notion of firing to Kahn process networks, as in [10, 7].

In the following we use projections from \mathbb{S}^m to \mathbb{S}^n , with $m \geq n$. These are functions that extract from an m -tuple an n -tuple such that for each value its number of occurrences in the argument is not smaller than the number of occurrences in the result. We write projections using the letter π with appropriate subscripts, and use them to map the input and output of an actor onto a subset of its ports. A special kind of projection maps a tuple onto one port, say p , of the actor. We write this projection as π_p .

10.2 Time systems

We want to allow actors to have a notion of time, without committing to any particular system of time stamps (or "tags", cf. [12]). Instead, we introduce the notion of time systems, which are basically sets of time stamps with the properties and operations we need to manipulate them.

Definition 5 (Time system). A *time system* is a structure

$$(T, \Delta, z, +, \leq)$$

with the following properties:

1. T and Δ are sets of time *tags* and *delays*, respectively.
2. \leq is a reflexive partial order on T
3. $+$: $T \times \Delta \rightarrow T$ is defined such that $\forall t \in T, \delta \in \Delta : t \leq t + \delta$
4. $z \in \Delta$ is the *zero delay* such that $\forall t \in T : t + z = t$.

Many time systems have a temporal metric associated with them that defines a quantitative measure of the "distance" between two time tags. If we use a time system together with such a metric we call it a metric time system.

temporal
metric

Definition 6 (Temporal metric/distance). Given a time system $(T, \Delta, z, +, \leq)$, a partial function

$$d : T \times T \rightarrow \mathbb{R}$$

is a *temporal metric* of this system iff it has the following properties for all $t_1, t_2, t_3 \in T$:

- $t_1 \leq t_2 \Rightarrow d(t_1, t_2) \geq 0$
- $d(t_1, t_2) = -d(t_2, t_1)$
- $t_1 \leq t_2 \leq t_3 \Rightarrow d(t_1, t_2) \leq d(t_1, t_3) \wedge d(t_2, t_3) \leq d(t_1, t_3)$
- $t_1 \leq t_2 \leq t_3 \Rightarrow d(t_1, t_2) + d(t_2, t_3) \geq d(t_1, t_3)$

For any two time tags t_1, t_2 the value $d(t_1, t_2)$, if defined, is called their (*directed*) *temporal distance*.

Note. d differs from what is usually called a “metric” in several important ways. First of all it is a partial function, and is required to only quantify distances between time tags that are comparable inside the time system. Also, it allows distinct time tags to have a zero temporal distance, and it reflects the “directedness” of the time system by allowing for negative distances.

Since a temporal metric is not required to return a non-zero value for distinct time tags, there is always a trivial temporal metric for each time system, viz. the one that yields 0 for any two tags.

temporal
“metric”
trivial
temporal
metric

Example 29. Some common time systems are the following:

- A familiar time system would be

$$(\mathbb{R}, \mathbb{R}_0^+, \leq, +)$$

where the comparisons and operators are the common ones on the real numbers. Its temporal metric is subtraction.

- A somewhat more sophisticated time system would be one that employs a notion of *delta time*:

$$(\mathbb{R} \times \mathbb{N}, \mathbb{R}_0^+ \times \mathbb{N}, \leq_{\text{LEX}}, +^n)$$

The time tags are tuples of a real time stamp and an additional natural number index. Here, $T = \mathbb{R} \times \mathbb{N}$, $\Delta = \mathbb{R}_0^+ \times \mathbb{N}$, $+$ adds two tuples positionwise, the order is the lexicographic order on the tuples, and the d is simply subtraction in the first (real-valued) component. For instance, $(-3.4, 2) + (5, 4) = (1.6, 6)$, and $(1.2, 50) \leq_{\text{LEX}} (1.21, 5) \leq_{\text{LEX}} (1.21, 6)$, and $(1, 6) \not\leq_{\text{LEX}} (2, 5)$.

A common metric that is used with this time system computes the difference of the first component, e.g. $d((1.2, 5), (1.21, 50)) = 0.01$.

- Another time system is *vector time*. Each time tag (and delay) is an n -ary vector of natural numbers:

$$(\mathbb{N}^n, \mathbb{N}^n, \leq^n, +^n)$$

Addition is the usual vector addition as in the previous example. The comparison compares two vectors positionwise, and two vectors are only comparable iff all their components are related to each other in the same way—e.g., $(1, 2, 3) \leq^3 (2, 3, 4)$ but $(1, 2, 3) \not\leq^3 (2, 3, 1)$ and also $(2, 3, 1) \not\leq^3 (1, 2, 3)$. In other words, unlike the previous two examples, this time system is not totally ordered.

It is also a system that is often not associated with a temporal metric. Many possible candidate metrics exist, however, such as sum or product of positionwise differences, or for each component of the vector the difference in that position.

10.3 Actor transition systems

We can now define an actor as an entity that has input/output ports through which it interacts with its environment by receiving and sending tokens, that makes atomic steps, and that in each such step consumes and produces a finite (and possibly empty) sequence of tokens at each input or output port. It also has a state (which we will not further characterize), which may change during this step.

Definition 7 (Actor transition system, actor, ATS). Let \mathcal{U} be the *universe* of all values, and $\mathbb{S} = \mathcal{U}^*$ be the set of all finite sequences in \mathcal{U} . Further, let $(T, \Delta, z, +, \leq)$ be a time system. For any non-empty set Σ of *states* an *m-to-n actor transition system* (or just *actor* or *ATS* for short) is a labeled transition system

$$\langle \sigma_0, \Delta, \tau, \succ \rangle$$

with $\sigma_0 \in \Sigma$ its *initial state*, and

$$\tau \subseteq \Sigma \times \mathbb{S}^m \times \Delta \times \mathbb{S}^n \times \Sigma$$

its *transition relation*. Any $(\sigma, s, \delta, s', \sigma') \in \tau$ is called a *transition*. σ is its *source state*, σ' its *destination state*, s its *input tuple*, s' its *output tuple*, and δ its *delay*. (s, δ, s') are the *transition label*.

Finally, \succ is a non-reflexive, anti-symmetric and transitive partial order relation on τ , called its *priority relation*.

Notation 1. For any transition $(\sigma, s, \delta, s', \sigma') \in \tau$ we also write

$$\sigma \xrightarrow[\delta]{s \mapsto s'} \sigma'$$

or, if δ is understood or not relevant,

$$\sigma \xrightarrow{s \mapsto s'} \sigma'$$

The set of all *m-to-n actors* with firing is $\mathcal{A}^{m \rightarrow n}$. The set of all actors is

$$\mathcal{A} =_{def} \bigcup_{m, n \in \mathbb{N}} \mathcal{A}^{m \rightarrow n}$$

The core of this definition can be found in [10], where “firing rules” defined the input tuples and a firing function mapped those to output tuples. State was added in an extension proposed in [7]. Here, we add the priority relation, which makes actors much more expressive, by allowing them to ascertain and react to the “absence” of tokens. On the other hand, it can make them harder to analyze, and it may introduce unwanted non-determinism into a dataflow model.

Intuitively, the priority relation determines that a transition cannot occur if some other transition is possible. We can see this in the definition of a valid

priority
relation
allows test for
absence of
tokens

step of an actor, which is a transition such that two conditions are satisfied: the required input tokens must be present, and there must not be another transition that has priority.

Definition 8 (Enabled transition, step). Given an m-to-n actor $\langle \sigma_0, \Sigma, \tau, \succ \rangle$, a state $\sigma \in \Sigma$ and an *input tuple* $v \in \mathbb{S}^m$, a transition $\sigma \xrightarrow[\tau]{s \mapsto s'} \sigma'$ is *enabled* iff

$$s \sqsubseteq v \\ \neg \exists \sigma \xrightarrow{r \mapsto r'} \sigma'' \in \tau : r \sqsubseteq v \wedge \sigma \xrightarrow{r \mapsto r'} \sigma'' \succ \sigma \xrightarrow{s \mapsto s'} \sigma'$$

A *step* from state σ with input v is any enabled transition $\sigma \xrightarrow{s \mapsto s'} \sigma'$. The *residual input tuple* v' is defined by $v = s + v'$.

Note that the second condition for a transition to be enabled becomes vacuously true if $\succ = \emptyset$, leaving $s \sqsubseteq v$, the usual dataflow condition. We call an actor with an empty priority relation a *pure actor*.

pure actor

Part III

Appendices

Appendix A

CAL language syntax

The following is a summary of the CAL language syntax. See chapter 2 for some relevant conventions and the syntax of lexical tokens.

A.1 Actor

```
Actor → [Imports] actor ID
      [ '[' TypePars ']' ] '(' ActorPars ')' IOSig [TimeClause] ':'
      { VarDecl | Action | InitializationAction | PriorityOrder }
      [ActionSchedule]
      { VarDecl | Action | InitializationAction | PriorityOrder }
      (end|endactor)

TypePar → ID [ '<' Type ]
ActorPar → [Type] ID [=] Expression
IOSig → [PortDecls] '==>' [PortDecls]
PortDecl → [multi] [Type] ID
TimeClause → time Type
Import → SingleImport | GroupImport ';'
SingleImport → import QualID [=] ID
GroupImport → import all QualID
QualID → ID { . ID }
Type → ID
      | ID '[' TypePars ']'
      | ID '(' [TypeAttr { ',' TypeAttr }] ')'
      | '[' [Types] '-->' Type ']'
      | '[' [Types] '-->' ']'
```

TypeAttr \rightarrow ID ':' Type
 | ID '=' Expression
 VarDecl \rightarrow [**mutable**] [Type] ID [('=' | ':=') Expression]
 | FunDecl | ProcDecl

A.2 Expressions

Expression \rightarrow PrimaryExpression { Operator PrimaryExpression }
 PrimaryExpression \rightarrow [Operator] SingleExpression
 { '(' [Expressions] ')' | '[' Expressions ']' | '.' ID }
 SingleExpression \rightarrow [**old**] ID
 | ExpressionLiteral
 | '(' Expressions ')'
 | IfExpression
 | LambdaExpression
 | ProcExpression
 | LetExpression
 | ListComprehension
 | SetComprehension
 | MapComprehension
 ExpressionLiteral \rightarrow IntegerLiteral | DecimalFractionLiteral
 | StringLiteral
 | **true** | **false** | **null**
 IfExpression \rightarrow **if** Expression **then** Expression **else** Expression **end**
 LetExpression \rightarrow **let** VarDecls ':' Expression (**end|endlet**)
 LambdaExpression \rightarrow [**const**] **lambda** '(' [FormalPars] ')' ['-->' Type]
 [**var** VarDecls] ':' Expression (**end|endlambda**)
 FormalPar \rightarrow [Type] ID
 ProcExpression \rightarrow **proc** '(' [FormalPars] ')' [**var** VarDecls]
 (**do|begin**) { Statement } (**end|endproc**)
 FuncDecl \rightarrow **function** ID '(' [FormalPars] ')' [**var** VarDecls ':'] Expression **end**
 ProcDecl \rightarrow **procedure** ID '(' [FormalPars] ')' [**var** VarDecls (**begin|do**)]
 { Statement } **end**

SetComprehension \rightarrow '{' [Expressions [':' Generators]] '}'
 ListComprehension \rightarrow '[' [Expressions [':' Generators] ['|' Expression]] '
 MapComprehension \rightarrow **map** '{' [Mappings [':' Generators]] '}'
 Mapping \rightarrow Expression '->' Expression
 Generator \rightarrow **for** [Type] ID [',' IDs] **in** Expression { ',' Expression }
 TypeAssertionExpr \rightarrow '(' Expression '::' Type ')'

A.3 Statements

Statement \rightarrow AssignmentStmt
 | CallStmt
 | BlockStmt
 | IfStmt
 | WhileStmt
 | ForeachStmt
 | ChooseStmt
 AssignmentStmt \rightarrow ID [Index | FieldRef] ':=' Expression ';'

Index \rightarrow '[' [Expressions] '
 FieldRef \rightarrow '.' ID
 CallStmt \rightarrow Expression '(' [Expressions] ')';'
 BlockStmt \rightarrow **begin** [var VarDecls] **do** { Statement } **end**
 IfStmt \rightarrow **if** Expression **then** { Statement } [**else** { Statement }] **end**

WhileStmt \rightarrow **while** Expression [var VarDecls] **do** [Statements] (**end|endwhile**)
 ForeachStmt \rightarrow ForeachGenerator { ',' ForeachGenerator }
 [var VarDecls] **do** [Statements] (**end|endforeach**)
 ForeachGenerator \rightarrow **foreach** [Type] ID { ',' ID } **in** Expression
 [',' Expressions]

ChooseStmt \rightarrow ChooseGenerator { ',' ChooseGenerator }
 [**var** VarDecls] **do** [Statements]
 [**else** [[**var** VarDecls] **do**] [Statements]] (**end|endchoose**)
 ChooseGenerator \rightarrow **choose** [Type] ID [',' IDs] **in** Expression [',' Expressions]

A.4 Actions

Action \rightarrow [ActionTag ':'] **action** ActionHead [**do** Statements] (**end|endaction**)
 ActionTag \rightarrow ID { ':' ID }
 ActionHead \rightarrow InputPatterns '==>' OutputExpressions
 [**guard** Expressions] [**var** VarDecls] [**delay** Expression]
 InputPattern \rightarrow [ID ':'] '[' IDs ']' [RepeatClause] [ChannelSelector]
 ChannelSelector \rightarrow **at** Expression
 | **at*** Expression
 | [**at***] **any**
 | [**at***] **all**
 RepeatClause \rightarrow **repeat** Expression
 OutputExpression \rightarrow [ID ':'] '[' Expressions ']' [RepeatClause]
 [ChannelSelector]
 InitializationAction \rightarrow [ActionTag ':']
 initialize InitializerHead [**do** Statements] (**end|endinitialize**)
 InitializerHead \rightarrow '==>' OutputExpressions
 [**guard** Expressions] [**var** VarDecls] [**delay** Expression]

Appendix B

Keywords

The following table lists the currently active reserved words in CAL. The comment either names the context of its use or provides an example schema of its use. Note that these are not complete syntax rules. The references point to further information on the keyword.

Keyword	Comment	Reference
action	action actionhead do statements end	8
actor	actor actorhead : decls, actions, etc. end	3
all	input patterns	8.1
	output expressions	8.2
and	Boolean operator	C.1, C.2.9
any	input patterns	8.1
	output expressions	8.2
at	input patterns	8.1
	output expressions	8.2
at*	input patterns	8.1
	output expressions	8.2
begin	procedure head begin statements end	6.9.3
	proc head begin statements end	6.9.2
	begin statements end	7.3
choose	choose generators do statements end	7.7
const	const lambda head : expression end	6.9.1
delay	action head	8.3
div	arithmetic operator	C.1, C.2.6
do	actions	8
	structured statements	7.5, 7.6, 7.7

continued on next page

KEYWORDS

		(continued)
Keyword	Comment	Reference
dom	domain operator on maps	C.1, C.2.5
else	if expression then expression else expression end	6.7
	if expression then statements else statements end	7.4
	choose generators do statements else statements end	7.7
end	generic end marker	
endaction	alternative end marker, see action	
endactor	alternative end marker, see actor	
endchoose	alternative end marker, see choose	
endforeach	alternative end marker, see foreach	
endfunction	alternative end marker, see function	
endif	alternative end marker, see if	
endinitialize	alternative end marker, see initialize	
endlambda	alternative end marker, see lambda	
endlet	alternative end marker, see let	
endpriority	alternative end marker, see priority	
endproc	alternative end marker, see proc	
endprocedure	alternative end marker, see procedure	
endschedule	alternative end marker, see schedule	
endwhile	alternative end marker, see while	
false	Boolean constant literal	C.2.9
for	for vars in expression	6.10.2
foreach	foreach generators do statements end	7.6
fsm	schedule fsm tag : transitions end	9.2.1
function	function head : expression end	6.9.3
guard	action header	8
if	if expression then expression else expression end	6.7
	if expression then statements else statements end	7.4
import	import qualid [= var]	3.1
in	membership operator, collections	C.1, C.2.1
	generators	6.10.2, 7.6, 7.7
initialize	initialize head do statements end	8.5
lambda	lambda head : expression end	6.9.1
let	let vardecls : expression end	6.8

continued on next page

KEYWORDS

(continued)

Keyword	Comment	Reference
map	map { expressions [: generators] }	6.10
mod	modulo operator	C.1, C.2.6
multi	multi-port declaration	3
mutable	mutable vardecl	5.1.1
not	Boolean operator	C.1, C.2.9
null	literal	2.1
old	old var	6.2.1
or	Boolean operator	C.1, C.2.9
priority	priority inequalities end	9.3
proc	proc head do statements end	6.9.2
procedure	procedure head do statements end	6.9.3
regexp	schedule regexp regexp end	9.2.2
repeat	input patterns	8.1
	output expressions	8.2
rng	range operator on maps	C.1, C.2.5
schedule	schedule schedule end	9.2
then	if expression then expression else expression end	6.7
	if expression then statements else statements end	7.4
time	actor head	3
true	Boolean constant literal	C.2.9
var	variable declarations marker	
while	while expression do statements end	7.5

The following keywords are reserved for future use.

```

assign    case    default    endinvariant    endtask
endtype   ensure  invariant   now              out
protocol  require  task       type

```

Appendix C

Basic runtime infrastructure

This appendix describes the basic runtime infrastructure, i.e. the kinds of objects and operations on them that implementations must provide in order to implement CAL. A realistic implementation might support many more data types and have a much more extensive library of functions and procedures. Most of the facilities described in this appendix are used in CAL constructs, such as collections in generators, or maps and sets in input patterns and output expressions.

C.1 Operator symbols

The following table summarizes the predefined unary operator symbols in CAL.

Operator	Operand type	Meaning
not	Boolean	logical negation
#	Collection[T]	number of elements
	Map[K, V]	number of mappings
dom	Map[K, V]	domain of a map
rng	Map[K, V]	range of a map
-	Number	arithmetic negation

The next table lists the predefined binary operator symbols in the CAL language. They are sorted by increasing binding strength. Their binding strength is given by a precedence figure P , higher precedence binds stronger.

P	Operator	Operand 1	Operand 2	Meaning
1	and	Boolean	Boolean	logical conjunction
	or	Boolean	Boolean	logical disjunction
2	=	Object	Object	equality
	!=	Object	Object	inequality
	<	Number	Number	less than
		Set[T]	Set[T]	
		String	String	
		Character	Character	
	<=	analogous to <		less then or equal
	>	analogous to <		greater than
>=	analogous to <		greater than or equal	
3	in	T	Collection[T]	membership
4	+	Number	Number	addition
		Set[T]	Set[T]	union
		List[T]	List[T]	concatenation
		Map[K, V]	Map[K, V]	map union
	-	Number	Number	difference
		Set[T]	Set[T]	set difference
5	div	Number	Number	integral division
	mod	Number	Number	modulo
	*	Number	Number	multiplication
		Set[T]	Set	intersection
	/	Number	Number	division
6	^	Number	Number	exponentiation

C.2 Basic data types and their operations

This section shortly lists the basic data types used by CAL, describes their fundamental properties, and the operations they can be expected to support.

Many of the basic data types are composite, i.e. elements of these types may be objects that contain other objects. In the cases presented in this section, such a composite can be recursively decomposed into a subobject and a residual composite, until a canonical empty composite is reached. Functions operating on such composites can often easily be described by using selector functions. A selector function takes a composite and another function as arguments. It picks one of the subobjects and passes it and the residual composite into the argument function. In the cases below, we can describe our composite data types by some properties of the data objects and the behavior of the selector function.

selector
function

C.2.1 Collection[T]—collections

An object of type `Collection[T]` is a finite collection of objects of type `T`. The most fundamental operation on it is selecting an element, which is really

dividing the collection into an element and a residual collection of the same kind. The `selectf` selector function does this, and has the following signature:

`selectf`

```
selectf<A, B>: [[A, Collection[A] --> B],
               Collection[A],
               B --> B]
```

Different kinds of collections may differ in the way the selector function picks the element from the collection (e.g., in the case of lists, it is guaranteed to pick the first element of the list, see [C.2.4](#)).

The two fundamental operations on a collection are computing its size and testing whether a given object is contained by it. These are represented by two operators, the prefix operator `#` and the infix operator `in`. We assume that these are substituted by function calls as follows:

`size and membership`

```
# C ≡ $size(C)
a in C ≡ $member(a, C)
```

We can define the `$size` and `$member` functions using the `selectf` function in the following manner:

```
function $size(C) : // prefix #
  selectf(lambda (a, C1) :
    1 + $size(C1)
  end, C, 0)
end
```

```
function $member(x, C) :
  selectf(lambda (a, C1) :
    if a = x then true else
      $member(x, C1)
    end
  end, C, false)
end
```

Finally, we can define iteration over a collection by writing a procedure that works somewhat analogous to the `selectf` function:

```
procedure selectp(p, C) begin
  selectf(lambda (a, C1) :
    proc () do p(a, C1); end
  end,
  C,
  proc () do end)();
end
```

Consequently, the signature of `selectp` is as follows:

```
selectp<A>: [[A, Collection[A] -->], Collection[A]]
```

Generator functions and procedures

The selection function and procedure also allow us to construct the various functions and procedures used in defining generators in CAL. For instance, the `$mapadd` function used in comprehensions (see section 6.10.2), which is slightly different depending on whether it produces a set, a list, or a map:

```
function $mapaddset (C, f) :
    selectf(lambda (a, C1) : f(a) + $mapaddset(C1, f) end,
           C, {})
end

function $mapaddlist (C, f) :
    selectf(lambda (a, C1) : f(a) + $mapaddlist(C1, f) end,
           C, [])
end

function $mapaddmap (M, f) :
    selectf(lambda (a, M1) : f(a) + $mapaddmap(M1, f) end,
           M, map {})
end
```

The `$iterate` procedure used in the `foreach`-statement (see section 7.6) can be built on top of `selectp` as follows:

```
procedure $iterate (C, p) begin
    selectp(proc (a, C1) do p(a); $iterate(C1, p); end, C);
end
```

Finally, the `$try` procedure used in the `choose`-statement (see section 7.7) can be constructed as follows:

```
procedure $try (C, f, p) begin
    selectp(proc (a, C1) do
        if f() then
            p(a);
            $try(C1, f, p);
        end
    end, C);
end
```

C.2.2 Seq[T]—sequences

Sequences are arrangements of objects that are indexed by non-negative integers, starting at zero if not empty, and either ending at some maximal index k , or not at all. The first kind of sequence is called finite sequence or list (see section C.2.4), the second kind is called infinite sequence.

The selector function of sequences splits the non-empty sequence into its first element and the rest of the sequence:

```
selectf<A, B>: [[A, Seq[A] --> B],
               Seq[A],
               B --> B]
```

The only operations on sequences are indexing (using a single integral index) and testing whether a given non-negative integer is a valid index. They can be expressed in terms of the selector as follows:

`hasElement`

```
function $nth(n, S) : // indexer
  selectf(lambda (a, R) :
    if n = 0 then a else
      $nth(n - 1, R)
    end, S, null)
end

function hasElement (n, S) :
  selectf(lambda (a, R) :
    if n = 0 then true else
      hasElement(n - 1, R)
    end, S, false)
end
```

C.2.3 Set[T] < Collection[T]—sets

Sets are a special kind of collection (see section C.2.1) that guarantee that each object occurs in them at most once. The selection function `selectf` is nondeterministic on sets with two or more elements, i.e. it is unspecified which element will be chosen.

`nondeterministic
selectf`

Sets are partially ordered by the subset order, which can be defined as follows:

`subset:
partial order`

```
function subset(S1, S2) :
  selectf(lambda (a, S) :
    if a in S2 then
      subset(S, S2)
    else
      false
    end
  end, S1, true)
end
```

The comparison operators (<, <=, >, >=) are based on the subset order. In addition, two sets are considered equal if they are subsets of each other, i.e. they contain the same elements.

`set equality`

The fundamental way of constructing sets is to add one element to an existing

`addElement`

set. This can be done using the function `addElement`, which has the following signature:

```
addElement<T> : [T, Set[T] --> Set[T]]
```

The result of `addElement(a, S)` is the smallest set that contains both, `a` and all the elements in `S`.

Sets also have operators for union (+), intersection (*) and set difference (-), which can be defined as follows:

```
function $union (S1, S2) : // infix +
  selectf(lambda(a, S) :
    addElement(a, $union(S, S2))
  end, S1, {})
end

function $intersection(S1, S2) : // infix *
  selectf(lambda(a, S) :
    if a in S2 then
      addElement(a, $intersection(S, S2))
    else
      $intersection(S, S2)
    end
  end, S1, {})
end

function $setDifference(S1, S2) : // infix -
  selectf(lambda(a, S) :
    if a in S2 then
      $setDifference(S, S2)
    else
      addElement(a, $setDifference(S, S2))
    end
  end, S1, {})
end
```

In addition to these basic operations, there are set comprehensions for computing sets (section 6.10).

C.2.4 List[T] < Collection[T], Seq[T]—lists

Lists are finite, sequential arrangements of objects. They are collections and sequences at the same time. As a consequence, its selector function splits the (non-empty) list into its first element and the rest of the list.

For any list, the value computed by the `$size` function (section C.2.1) defines the valid indices (via `$nth` and `hasElement`, see section C.2.2) into the list, which are the non-negative integers smaller than that number.

Similar to sets, lists can be constructed elementwise using a `cons` function, that adds an element to the front of an existing list:

`cons`
constructor

```
cons<T> : [T, List[T] --> List[T]]
```

Lists can be concatenated using the `+` operator, which can be defined as follows:

```
function $add (L1, L2) : // operator +
  selectf(lambda (a, L) :
    cons(a, $add(L, L2))
  end, L1, L2)
end
```

Similarly to sets, list can be created using list comprehensions (section 6.10).

C.2.5 Map[K, V]—maps

A map of type `Map[K, V]` maps finitely many keys of type `K` to values of type `V`. The set of keys of a map is called its domain, while the set of values is called its range.

Like collections, maps are accessed using a selector function. In the case of maps, this has the following signature:

```
selectf<K, V>: [[K, V, Map[K, V] --> A],
               Map[K, V],
               A --> A]
```

It separates a non-empty map into a key, the corresponding value, and the residual map, which it then applies its first argument to.

The domain and range of a map can then be defined like this:

```
function $domain(M) : // operator dom
  selectf(lambda (k, v, M1) :
    addElement(k, $domain(M1))
  end, M, {})
end

function $range(M) : // operator rng
  selectf(lambda (k, v, M1) :
    addElement(v, $range(M1))
  end, M, {})
end
```

The value corresponding to a key in a map is retrieved via an indexer, using the key as in index into the map. The indexing function is defined as follows:

```
function $get(x, M) // indexer
  selectf(lambda(k, v, M1) :
```



```

        if x = k then v else
            $get(x, M1)
        end
    end, M, null)
end

```

C.2.6 Number—numbers

Most of the number system of CAL is left to platform definition. The language itself provides two kinds of numeric literals (integers are decimal fractions). There are a number of operators on numbers, cf. section C.1 for details.

The type `Number` contains the subtype `Integer` of integer numbers.

C.2.7 Character—characters

Characters are usually related to other data types, such as numbers (via some encoding) or strings (as components of strings). The details of this are left to the definition of platforms.

Characters are (possibly partially) ordered, and thus support the relational operators `<`, `<=`, `>`, `>=`.

C.2.8 String < List[Character]—strings

Strings are lists of characters, supported by a special syntax for their construction. It is left to an implementation whether they are identical to a corresponding list of characters, i.e. whether

```
"abc" = ['a', 'b', 'c']
```

Strings ‘inherit’ the order from characters, and expand it lexicographically:

```

function $stringLE (s1, s2) : // operator <=
    selectf(lambda (c1, r1) :
        selectf(lambda (c2, r2) :
            if c1 = c2 then
                $stringLE(r1, r2)
            else
                c1 < c2
            end
        end, s2, false)
    end, s1, true)
end

```

C.2.9 Boolean—truth values

The `Boolean` data type represents truth values. It is supported by two constant literals (**true** and **false**) as well as a number of operators, cf. section C.1.

C.2.10 Null—the null value

The `Null` data type has only one representative, the object denoted by the literal **`null`**. Apart from equality tests, no operations are defined on this type.

Acknowledgments

Many people have contributed to this work in various ways, and the authors would like to acknowledge their contributions and thank them for their input and support.

Special thanks go to Edward A. Lee, Christopher Hylands, H. John Reekie, Anders Blomdell, Lars Wernli, Chris Chang, Yang Zhao, Ernesto Wandeler, Ed Willink, Steven Neuendorffer, Kees Vissers, C. Xavier Brooks, and the Ptolemy Group at UC Berkeley.

The research reported here is part of the Ptolemy project, which is supported by the Defense Advanced Research Projects Agency (DARPA), the MARCO/DARPA Gigascale Silicon Research Center (GSRC), the State of California MICRO program, and the following companies: Agilent Technologies, Cadence Design Systems, Hitachi, and Philips.

Part of this research was supported by STINT—The Swedish Foundation for International Cooperation in Research and Higher Education.

Addresses & contact

Johan Eker	johan.eker@ericsson.com
Jörn W. Janneck	jwj@acm.org
The Ptolemy Project	http://ptolemy.eecs.berkeley.edu
CAL Development	http://www.caltrop.org

Index

- (...)
 - in action schedule, 70, 71
 - in expression, 35
- *
 - in action schedule, 71
- >
 - in lambda-expression, 40
- >, 43
- - in action tag, 57
 - in assignments, 48
 - in expression, 35
 - in field selection expression, 38
- :
 - after action tag, 57
 - in action schedule, 70
 - in actor head, 17
 - in initialization action, 66
 - in input pattern, 58
 - in lambda-expression, 40
 - in output expression, 62
- ::, 47
- ::=
 - in assignments, 48
 - in variable declarations, 29
- ;
 - in action schedule, 70
 - in assignments, 48
 - in priority clause, 73
 - in procedure call, 51
- <
 - in type parameter constraint, 17
- =
 - in actor parameters, 17
 - in variable declarations, 29
- ==>
 - in action head, 57
 - in initialization action, 66
 - in port declaration, 17
- >
 - in priority clause, 73
- [...], 43
 - in action schedule, 71
 - in assignments, 48
 - in expression, 35
 - in indexing expression, 39
 - in input pattern, 58
 - in output expression, 62
- #
 - on collection, 92
- \$
 - in reserved identifiers, 13
- \$iterate
 - definition, 93
 - use in foreach-statement, 54
- \$mapadd
 - definition, 93
 - use in generators, 45
- \$try
 - definition, 93
 - use in choose-statement, 55
- {... }, 43
 - |
 - in action schedule, 70, 71
- action, 17, 57–67
 - activated, 65
 - activation conditions, 65
 - and transition, 57
 - body, 57

- choice, 57
- control structures, 68–74
- delay, 57, 65
- duration, 19
- firability, 58
- head, 57
- initialization action, 66–67
 - activation conditions, 66
 - multiple activated, 66
- legal sequence of, 70
- priority, 73–74
 - order construction, 73
 - partial order, 73
- schedule, 69–73
 - finite state machine, 70–71
 - regular expression, 71–73
- scope, 61–62
- selection, 65–66, 68–74
- tag, 57, 69
 - denoted actions, 69
 - extension, 69
 - language, 70, 71
 - prefix order, 69
- untagged
 - and schedules, 70
- variable
 - scope, 61–62
- action**, 57
- action schedule, 17, 69
- activation conditions
 - of action, 65
 - of initialization action, 66
- actor, 76–80
 - firing, 6
 - initial state, 79
 - parameter, 17, 18, 28
 - port signature, 17, 18
 - priority relation, 79
 - pure, 80
 - state, 6, 79
 - change, 57
 - structure of, 17
 - transition, 79
 - relation to action, 57
 - transition relation, 79
 - type parameter, 17
 - bound, 17
- actor**, 17
- actor model, 76–80
- actor transition system, 79
- actors
 - composition of, 6
 - scheduling of, 7
- aliasing, 49
- all**, 58, 60, 62
- any**, 58, 60, 62
- application expression, 37–38, 41
- art of cheating, 37
- assignability, see types, assignability
- assignment, see variable, 48–51
 - field, 48–49
 - indexed, 48–49
 - simple, 48–49
- associativity
 - of operators, 39
- at**, 58, 60, 62
- at***, 58, 60, 62
- ATS, see actor transition system
- backtracking
 - and choose-statement, 56
- Backus-Naur-Form, 15
- begin**, 41
- binary operator
 - predefined, 90
- binding, 28
 - assignable, 28
 - mutable, 28
- BNF, 15
- body
 - of expression, 16
 - of statement, 16
- Boolean (type), 24, 36, 40, 44, 52, 55, 97
- cast, see type assertion
- channel, 18
- channel identifier, 59
- channel selector, 58, 60
 - free variables of, 62
 - in output expression, 63
 - multichannel, 60, 64

- output expression
 - constraints on map domains, 64
 - single-channel, 60, 63
- ChannelID (type), 24
- Character (type), 97
- choose**, 54
- choose-statement, 54–56
 - backtracking, 56
 - filter, 54
 - semantics, 54
- cloning, 50
 - lazy, 50
- closure, 40–42
 - and **old** variables, 37
 - function, 38, 40–41
 - application, 41
 - application of, 40
 - const**, 41
 - invariant, 41
 - type, 41
 - parameter, 28
 - procedure, 40–42
 - call of, 40
 - relation to variable dependency, 30
- Collection (type), 24, 44, 91–93
 - membership test, 92
 - selector function, 92
 - size, 92
 - vs collection, 42
- collection
 - expression, 43
 - proper, 42
 - vs Collection, 42
- comment, 14
 - syntax, 14
- complete partial order
 - of sequences, 76
- composite object, 25, 39
- comprehension, 42–46
 - efficiency, 46
 - filter, 44
 - generator, 42, 44
 - element order, 44
 - semantics, 45
 - syntax, 43
 - implementation, 46
 - list, 43
 - tail-free, 45
 - map, 43
 - set, 43
 - simple, 43
 - with generators, 43–46
- condition
 - for firing, 57
- conditional expression, 39–40
- const**, 40, 41
- control flow, 52, 53
- control structures
 - action-level, 68–74
- corollary
 - mutual dependencies, 31
- cpo, see complete partial order
- dataflow, 7
- declaration
 - of function, 29, 42
 - of procedure, 29, 42
- declaration set
 - well-formed, 31
- definition
 - actor, 79
 - actor transition system, 79
 - ATS, 79
 - dependency relation, 32
 - dependency set, 30
 - enabled transition, 80
 - immediate dependency set, 30
 - legal action sequence, 70
 - step (of actor), 80
 - temporal metric, 77
 - time system, 77
 - well-formed declaration set, 31
- delay, 19–20, 57
 - in time system, 19
 - zero, 19
 - of action, 65
 - time, 77
- delay**, 57, 65, 66
- delimiter, 13
- dependency
 - circular, 30
- dependency relation, 32

- is partial order, 32
- dependency set, 30–34
 - definition, 30
 - example, 32
 - immediate, 30
- do**, 41, 52–54, 57, 66
- domain
 - of map, 96
 - definition, 96
- duration
 - of action firing, 19
- eager evaluation, 37
- else**, 39, 52, 54
 - in foreach-statement, 54
- end**, 17, 39–42, 52–54, 57, 66, 70, 71, 73
 - alternatives, 16
 - exception, 52
- end marker, 16
- endaction**, 57
- endactor**, 17
- endchoose**, 54
- endforeach**, 53
- endfunction**, 42
- endif**, 39, 52
- endinitialize**, 66
- endlambda**, 40
- endlet**, 40
- endpriority**, 73
- endproc**, 41
- endprocedure**, 42
- endschedule**, 70, 71
- endwhile**, 52
- environment, 40
 - global, 17, 18
- example
 - old** reference, 36
 - action variable scope, 62
 - actor
 - Basic FairMerge, 68
 - BiasedMerge, 74
 - FairMerge with priorities, 74
 - FairMerge with regexp schedule, 72
 - Nondeterministic Merge, 68
 - Symmetric FairMerge, 73
 - backtracking using choose-statement, 56
 - collections with generators, 44
 - FairMerge, with FSM schedule, 71
 - fields and indices, 25
 - foreach-statement, 53
 - \$mapadd function, 46
 - multichannel input patterns, 61
 - mutually recursive variable declarations, 32–33
 - non-numeric indices, 25
 - non-well-formed variable declarations, 33–34
 - operator associativity and precedence, 39
 - order in generator collections, 44
 - simple map expressions, 43
 - simple set and list expressions, 43
 - single-channel output expression, 63
 - single-port input pattern, 59
 - single-port input pattern with repeat-clause, 60
 - time systems, 78
- expression, 35–47
 - application, 40, 41
 - body, 16
 - closure, 40–42
 - const** function, 41
 - function, 40–41
 - invariant function, 41
 - procedure, 40–42
 - comprehension, 42–46
 - filter, 44
 - generator, 44
 - generator semantics, 45
 - list, 43
 - map, 43
 - set, 43
 - simple, 43
 - with generators, 43–46
 - conditional, 39–40
 - constant, 36
 - field selector, 38–39
 - function

- application, 37–38, 41
 - head, 16
 - indexing, 39
 - lambda, 38
 - lambda**, 40–41
 - const**, 41
 - list comprehension
 - tail-free, 45
 - literal, 36
 - decimal fraction, 36
 - integer, 36
 - string, 36
 - syntax, 36
 - local scope, 40
 - operator, 39
 - primary, 35
 - proc**, 41–42
 - selector, 38–39
 - side-effect-free, 35, 41
 - syntax, 35
 - type assertion, 46–47
 - type of, 35
 - variable, 28–34, 36–37
- extension
 - of action tag, 69
- false**, 24, 36, 55, 97
- field, 25
 - selector expression, 38–39
- field reference
 - in assignments, 48
- filter, 44
- firing
 - duration, 19
- firing condition, 57
- firing rules (from [10]), 79
- for**, 43, 44
- foreach**, 53
 - statement, 48
- foreach-statement, 53–54
 - semantics, 53
- free variable, 30, 33
- fsm**, 70
- function, 40–41
 - application, 37–38, 40, 41
 - const**, 41
 - declaration, 29, 42
 - invariant, 41
 - recursive, 30
 - type, see types, closure, 41
- function**, 42
- generator, 42–46
 - collection
 - order, 44
 - filter, 44
 - in foreach-statement, 53
 - semantics, 45
 - syntax, 43
 - variable, 44
 - scoping, 44
- generic type, see types, parametric
- group import, 19
- guard**, 57, 66
- hasElement
 - definition, 94
- head
 - of expression, 16
 - of statement, 16
- homogeneous use
 - of input pattern, 61
 - of output pattern, 64
- host environment, 14
- host language, 14
- identifier, 13
 - escape syntax, 13
 - escaped, rationale, see rationale,
 - escaped identifier
 - qualified, 18, 69
 - reserved, 13
- if**, 39, 52
 - statement, 48
- if-statement, 52
- implementation note
 - actor parameter, 18
 - cloning, 50
 - generator implementations, 46
 - lazy evaluation, 37
 - object types, 22

- output token activation condition, 66
- import**, 18–19
- importing
 - group import, 19
 - of binding, 18
 - single import, 19
- in**, 43, 44, 53, 54
- in (operator)
 - on collection, 92
- index, 25
 - in assignments, 48
- indexing expression, 39
- initial state
 - of actor, 79
- initialization action, 66–67
 - activation conditions, 66
 - multiple activated, 66
- initialization expression, 29
 - of a variable, 30
- initialization rule, 17
- initialize**, 66
- input pattern, 58–62
 - association with port, 58
 - channel selector, 58
 - homogeneous use, 61
 - implicit variable declaration, 58
 - multichannel, 59
 - multiport, 60–61
 - single-channel, 59
 - single-port, 59–60
 - variable
 - dependencies, 62
 - type, table, Fig. 8.1, 59
- input token, 57
- Integer (type), 25, 97
- key
 - of map, 96
- keyword construct, 15
- keywords, 13
- Kleene operator
 - in action schedule, 71
- labeled transition system, 79
- lambda**, 40
 - closure
 - type, see types, closure
 - expression, 38, 40–41
 - const**, 41
- language
 - core, 8
 - domain-specific, 8
 - usability, 8
- language design
 - goals, 7–9
- lazy evaluation, 37
- let**, 29, 40
- lexical scoping, 30
- lexical token, 13
- lexicographical order, 97
- List (type), 24, 95–96
 - concatenation, 96
 - cons, 95
 - constructor, 95
 - hasElement, 95
 - selector function, 95
 - size, 95
 - valid indices, 95
- list
 - comprehension, 42, 43
 - simple, 43
 - tail-free, 45
 - construction, 42
- literal, 36
 - decimal fraction, 36
 - integer, 36
 - numeric, 14
 - string, 36
- local name, 18
- location, 25–26, 39
 - field, 25
 - independence of, 26
 - index, 25
 - orthogonality, 26
- LTS, see labeled transition system
- Map (type), 24, 61, 64, 96–97
 - domain, 96
 - definition, 96
 - indexer, 96
 - key, 96

- range, 96
 - definition, 96
- selector function, 96
- value, 96
- map
 - comprehension, 42, 43
 - simple, 43
 - construction, 42
 - nondeterminism, 43
 - key, 43
 - value, 43
- map**, 43
- membership test
 - in collection, 92
- metric
 - of time system, see temporal metric
- metric time system, 77
- model, 6
 - communication, 7
- model of computation, 7
- multi**, 17
- multichannel channel selector, 60, 64
- multichannel input pattern, 59
- multichannel output expression, 64
- multiport, see port, multiport, 59
 - and input pattern, 58
- multiport input pattern, 60–61
- mutable**, 29
- mutation, 48
- \mathbb{N} (natural numbers), 76
- namespace, 18–19
 - hierarchical, 17, 18
 - local name, 18
 - subnamespace, 18
- nondeterminism
 - choose-statement, 54–56
- notation
 - $\sigma \xrightarrow[\delta]{s \mapsto s'} \sigma'$, 79
- notational idiom, 15
- Null (type), 24, 36, 98
- null**, 24, 36
- Number (type), 97
- numeric literal, see literal, numeric
- object
 - mutable, 26
 - structured, 25–26
- object modification, see types, mutable
- object mutation, see types, mutable
- object types, see types, object
- old**, 35–37
 - and closures, 37
 - meaning, 37
 - restrictions, 37
 - translation, 37
 - use in action, 36
- operator, 13, 39
 - associativity, 39
 - basic, 90–91
 - binary, 90
 - unary, 90
 - binary, 39
 - binding strength, 90
 - infix, 39
 - precedence, 39, 90
 - predefined, 90–91
 - binary, 90
 - unary, 90
 - prefix, 39
 - symbols, 90
 - unary, 39
- operators
 - represented by functions, 39
- order
 - lexicographical, 97
 - of time tags, 19
- output expression, 62–65
 - association with port, 58
 - multichannel, 64
 - single-channel, 63
 - token expression
 - type, table (Fig. 8.2), 63
- output pattern
 - homogeneous use, 64
- output token, 57
- package, see subnamespace
- parameter
 - actor, 28
 - closure, 28

- generic type, 17
 - bound, 17
 - of actor, 17, 18
- partial order
 - dependency relation is, 32
- pattern variable
 - dependencies, 62
 - scoping, 61–62
 - type
 - table, Fig. 8.1, 59
- platform
 - profile, 10
- platform independence, 9–10
- port
 - and input patterns, 58
 - and output expression, 58
 - input, 6
 - multiport, 18, 59
 - and input pattern, 58
 - output, 6
 - single, 18
 - single port, 59
 - and input pattern, 58
- port signature, see actor, port signature
- portability, see platform independence
 - language, 10
 - source code, 10
 - target code, 10
- precedence
 - of operators, 39
- prefix
 - of input, 6
- prefix order, 76
 - is complete partial order, 76
- prefix-monotonicity, 73
- primary expression, 35
- priority
 - of action, 73–74
 - order construction, 73
 - partial order, 73
 - order, 73
 - construction, 73
 - validity, 73
- priority**, 73
- priority block, 17
- priority relation
 - of actor, 79
- proc**, 41
- proc expression, 41–42
- procedural closure
 - type, see types, closure
- procedure, 40–42
 - call, 40, 51–52
 - declaration, 29, 42
- procedure**, 42
- profile, see platform, profile
- Ptolemy, 6, 7
- pure actor, 80
- qualified identifier, see identifier, qualified, 69
- range
 - of map, 96
 - definition, 96
- rationale
 - choose-statement, 55
 - escaped identifier, 14
 - literals, 36
 - \$mapadd function, 45
 - optional types, 21
 - output expression design, 64
 - redundant action schedule syntax, 72
 - variable declaration
 - well-formedness, 32
- recursion, see recursive closure
- recursive closure, 30
- recursive function, 30
- regexp, see regular expression
- regexp**, 71
- regular expression
 - in action schedule, 71
- repeat**, 58, 62
- repeat-clause, 63, 64
 - expression
 - free variables of, 62
 - in input pattern, 59
 - in output expression, 63
- § (finite sequences)

- $= U^*$, 76
- \mathbb{S}_∞ (sequences)
 - $\mathbb{S} \cup \mathbb{S}^{\mathbb{N}}$, 76
 - complete partial order, 76
 - prefix order, 76
- schedule
 - of actions, 70–73
 - finite state machine, 70–71
 - regular expression, 71–73
- schedule**, 70, 71
- scope
 - of variable, see variable, scoping
- scoping
 - lexical, 30
- selectf**
 - for collection, 92
 - for list, 95
 - for map, 96
 - for sequence, 93
 - for set, 94
- selector
 - expression, 38–39
- selector function, 91
 - of collection, 92
 - of list, 95
 - of map, 96
 - of sequence, 93
 - of set, 94
- selector procedure
 - of collection, 92
- selectp**
 - definition, 92
- Seq** (type), 24, 93–94
 - finite, see **List** (type)
 - hasElement**
 - definition, 94
 - index into, 93
 - indexer
 - definition, 94
 - infinite, 93
 - selector function, 93
- sequence
 - is complete partial order, 76
 - prefix order, 76
- Set** (type), 24, 94–95
 - addElement**, 94
 - comparison operators, 94
 - construction, 94
 - difference, 95
 - equality, 94
 - intersection, 95
 - selector function, 94
 - subset order, 94
 - union, 95
- set
 - comprehension, 42, 43
 - simple, 43
 - construction, 42
- side-effect free, 41
- single import, 19
- single port, 59
 - and input pattern, 58
- single-channel channel selector, 60, 63
- single-channel input pattern, 59
- single-channel output expression, 63
- single-port input pattern, 59–60
- size
 - of collection, 92
- state
 - of actor, 79
 - change, 57
- statement, 48–56
 - assignment, 48–51
 - field, 48–49
 - indexed, 48–49
 - simple, 48–49
 - begin-end, 52
 - block, 48, 52
 - body, 16
 - call, 40, 51–52
 - choice, 48
 - choose**, 54–56
 - backtracking, 56
 - semantics, 54
 - foreach**, 48, 53–54
 - generator, 53
 - semantics, 53
 - head, 16
 - if**, 48, 52
 - list of, 41
 - procedure call, 48, 51–52
 - while**, 48, 52–53

- String (type), 25, 97
 - order, 97
- subnamespace, 18
- subset order, 94
- substitutability, see types, substitutable
- subtype relation, 22, 27
- tag
 - in time system, 19
 - order, 19
 - of action, 69
 - denoted actions, 69
 - extension, 69
 - language, 70, 71
 - prefix order, 69
 - time, 77
- temporal distance, see temporal metric
- temporal metric, 77
 - definition, 77
 - trivial, 78
 - vs proper metric, 77
- then**, 39, 52
- time, 19–20
 - adding delays, 19
 - addition, 77
 - delay, 19, 77
 - zero, 19
 - order, 77
 - tag, 19, 77
 - order, 19
 - zero delay, 77
- time**, 17, 20
- time system, 19–20, 77–78
 - definition, 77
 - metric, 77
- token, 6, 76
 - input, 57
 - lexical, see lexical token
 - output, 57
 - test for absence of, 73, 79
- token expression
 - table (Fig. 8.2), 63
 - type of, 63
- transition
 - and action, 57
 - of actor, 79
- transition relation
 - of actor, 79
- transition system, see actor transition system
- true**, 24, 36, 40, 44, 52, 97
- type assertion, 46–47
- type bound, see actor, type parameter, bound
- type cast, see type assertion
- type constructor, 23
- type framework, 27
 - finite cover, 27
 - least upper bound, 27
 - subtype order, 27
- type inference, 46
- type parameter, 23
 - value, 23
- types, 21–27
 - assignability, 22, 27
 - closure, 23
 - function, 41
 - free, 26
 - function, 41
 - mutable, 26, 28, 48
 - object, 21–23
 - of token, 59
 - optional, 21
 - parametric, 23
 - required, 24
 - simple, 23
 - structured, 25–26, 39, 48
 - substitutable, 22
 - subtypes, see subtype relation
 - syntax, 23–24
 - variable, 21–23
- \mathcal{U} (the universe of values), 76
- unary operator
 - predefined, 90
- universe
 - of token values, 76
- value
 - of map, 96
- var**, 29, 40–42, 52–54, 57, 66
- variable, 28–34, see assignment

- assignable, 49
 - binding, 28
 - assignable, 28
 - mutable, 28
 - bound, 28
 - declaration, 28
 - actor parameter, 28
 - closure parameter, 28
 - explicit, 28–29
 - in function, 40
 - in procedure, 41
 - input pattern, 28
 - not well-formed, 30
 - declaration set
 - well-formed, 31
 - dependency, 30
 - circular, 30
 - dependency relation, 32
 - dependency set, 30–34
 - definition, 30
 - example, 32
 - immediate, 30
 - free, 18, 30, 33
 - generator, 44
 - initialization expression, 29, 30
 - mutable, 26
 - name, 29
 - old**, 36–37
 - and closures, 37
 - meaning, 37
 - restrictions, 37
 - translation, 37
 - use in action, 36
 - reference, 36–37
 - scoping, 30–34
 - in actions, 61–62
 - lexical, 30
 - shadowing, 30
 - stateful, 29, 41
 - stateless, 29
 - type, 29
- variable type, see types, variable
- while**, 52
 - statement, 48
- while-statement, 52–53

Bibliography

- [1] The Ptolemy Project. Department EECS, University of California at Berkeley (<http://ptolemy.eecs.berkeley.edu>). 6
- [2] Harold Abelson and Gerald Jay Sussman. Structure and Interpretation of Computer Programs. MIT Press, 2nd edition, 1999. 5
- [3] Gul A. Agha. Actors: A Model of Concurrent Computation in Distributed Systems. The MIT Press Series in Artificial Intelligence. MIT Press, 1986. 6
- [4] Gul A. Agha, Ian A. Mason, Scott F. Smith, and Carolyn L. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 1993. 6
- [5] J. B. Dennis. First version data flow procedure language. Technical Memo MAC TM 61, MIT Lab. Comp. Sci., May 1975. 6
- [6] Carl Hewitt. Viewing control structures as patterns of passing messages. *Journal of Artificial Intelligence*, 8(3):323–363, June 1977. 6
- [7] Jörn W. Janneck. Syntax and Semantics of Graphs—An approach to the specification of visual notations for discrete event systems. PhD thesis, ETH Zurich, Computer Engineering and Networks Laboratory, July 2000. 6, 77, 79
- [8] Jörn W. Janneck. Actors and their composition. Technical Report UCB/ERL 02/37, University of California at Berkeley, 2002. 7
- [9] Gilles Kahn. The semantics of a simple language for parallel programming. In *Proceedings of the IFIP Congress*. North-Holland Publishing Co., 1974. 6
- [10] Edward A. Lee. A denotational semantics for dataflow with firing. Technical Report UCB/ERL M97/3, EECS, University of California at Berkeley, January 1997. 6, 73, 77, 79, 104
- [11] Edward A. Lee. Embedded software. In M. Zelkowitz, editor, *Advances in Computers*, volume 56. Academic Press, 2002. to appear. 7

- [12] Edward A. Lee and Alberto Sangiovanni-Vincentelli. A denotational framework for comparing models of computation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17(12):1217–1229, December 1998. [7](#), [77](#)
- [13] Charles S. Pierce. How to make our ideas clear. In P. P. Wiener, editor, *Values in a Universe of Chance*. 1878. [5](#)