

EBBA Modelling Tool
a.k.a
Event Definition Language

Peter Bates

COINS Technical Report 87-35
April 1987 (revised: October 30, 1987)

Laboratory for Distributed Computing
Computer and Information Science Department
University of Massachusetts
Amherst, Massachusetts 01003

This research supported in part by the National Science Foundation under grants MCS-8306327, DCR-8318776, and DCR-8500332. And by the Defense Advanced Research Projects Agency, monitored by the Office of Naval Research under contract NR049-041.

CONTENTS **2**

Contents

1. Introduction	1
2. Lexical Elements of EDL Descriptions	2
2.1 Identifiers and Keywords	2
2.2 Constant Values	3
2.2.1 Numbers	3
2.2.2 Time values	4
2.2.3 Strings	4
3. Directives to the Model builder	5
3.1 File Inclusion	5
3.2 Library Switching	5
3.3 Creating Background Libraries—Merge	6
4. Describing High-Level Event Models	6
4.1 Event Heading	6
4.2 Event Expression	7
4.2.1 Sequencing	8
4.2.2 Repetition	9
4.2.3 Choice	10
4.2.4 Concurrency	10
4.3 Cond Clauses	12
4.3.1 Event Heading Parameters	13
4.3.2 Event Attributes	14
4.3.3 Event Expression “lexical” Functions	15
4.3.4 Environment Specific Functions	15
4.4 With Clauses	16
5. Describing Primitive Events	17
6. A Complete Example	18
A. Using the Model builder	21
B. Model Builder outputs	22
B.1 Diagnostic and informative messages	22
B.2 Real Compiler Outputs	24
C. BNF Description of EDL	26

**EBBA Modelling Tool
a.k.a
Event Definition Language**

Peter Bates

April 1987 (revised: October 30, 1987)

Laboratory for Distributed Computing
Computer and Information Science Department
University of Massachusetts
Amherst, Massachusetts 01003

ABSTRACT

This report describes how to use the Model Builder and the Event Definition Language (*EDL*) as implemented in the *EBBA* toolset. The *EBBA* toolset is a collection of programs that allow its users to monitor (and experiment with) a program distributed over a network of processors. The Event Definition Language (*EDL*) is a mechanism that enables its users to express models of system behavior as abstractions of the detailed computation comprising that behavior. *EDL* users express their behavioral abstractions as high-level event clusters in terms of the primitive events that characterize the system, and high-level events that were previously defined through that viewpoint.

This report describes a model building tool based on the *EDL* that is integrated with the other components of the *EBBA* toolset. The model builder is largely a translator for *EDL*-described behavior models. However, the language that it accepts is a bit richer than *EDL* in that it accepts directives that control compilation and event library management.

1. Introduction

The EBBA toolset is a collection of programs that allow its users to monitor (and experiment with) a program distributed over a network of processors. The philosophy and implications of this high-level debugging paradigm are developed in a thesis¹ and a detailed description of the toolset internals is found in another report². This report describes how to use the Model Builder and the Event Definition Language (*EDL*) as implemented in the *EBBA* toolset³.

The Event Definition Language is a mechanism that enables its users to express models of system behavior as abstractions of the detailed computation comprising that behavior. *EDL* users express their behavioral abstractions as high-level event clusters in terms of the primitive events that characterize the system, and high-level events that were previously defined through that same viewpoint. Behavior models expressed as *EDL* definitions themselves define event tuples which may be incorporated into other high-level event descriptions. Thus, each *EDL* event definition provides a template for an event tuple that represents an event class which might occur in the system.

The model builder is basically a translator for *EDL* described behavior models. The language that it accepts is a bit richer than *EDL* in that it accepts directives that control compilation and event library management. The model builder can be executed interactively through the model builder interface or in a simple batch mode on collections of files. In its interactive mode, the model builder can be linked with graphic display tools that assist users in visualizing the structures of their higher level models.

Both primitive and high-level events may be described in *EDL*. Primitive events result from system operation so users of *EBBA* descriptive tools have little control over creation of instances of primitive event classes. Hence, primitive events require little descriptive power. High-level events on the other hand, represent user models of behavior and in a sense are a creation of the *EBBA* model matching tools. They require more powerful means to describe their resultant event tuples as well as means for describing the intended behavioral model.

Primitive events described in *EDL* simply indicate their unique system-wide event names and supply a list of names for the attributes possessed by the event. Descriptions of higher level events are more elaborate since they involve use of *EBBA* clustering and filtering techniques. Clustering to describe a high-level model is effected by specifying a set of previously defined events in an event expression and imposing constraints on the event expression members' attributes in a set of relational expressions. The event and relational expressions serve in dual roles of describing a

¹P.C. Bates, "Debugging Programs in a Distributed System Environment", COINS Technical Report 86-05, January 1986

²P.C. Bates, "The EBBA Toolset Implementation Notes"

³P.C. Bates, "Distributed Debugging tools for Distributed Systems", COINS Tech Report 87-28, March 1987

behavioral model and providing a guide for recognizing the occurrence of an event representing the model. The attributes of a high-level event are expressed in terms of the attributes defined by the cluster member events and other information obtainable from the *EBBA* model matching environment.

The next two sections describe the lexical elements from which *EDL* definitions are composed and the form of the input language to the model builder. The following section explains in detail how high-level events are written in *EDL*. Section four describes how the primitive events which characterize a system viewpoint are described in *EDL* and relates these event definitions to the event templates moved around the system. The final section provides an example of the use of *EDL* to describe a couple of behavior models.

2. Lexical Elements of *EDL* Descriptions

EDL event descriptions are composed from a limited set of lexical elements including keywords, identifiers, values, expression operators, and punctuation marks. Keywords introduce clauses and special values and will be written here in bold face letters, e.g. **event**, **end**, **primitive**, **with**, **cond**, **true**. Use of individual keywords will be described in appropriate sections.

2.1 Identifiers and Keywords

Identifiers are constructed from contiguous sequences of letters, digits, and underscores with the restriction that the first be a letter. Identifiers are used to represent event attributes, event class names and parameters, constants, and external functions. Identifiers will be written in italic throughout this and following sections. Some valid identifiers would be

Send *CreateHyp* *Create.Special.Node*
Event3 *TIME_OF_DAY* *legitimate.10.folds*

Identifier strings are not case sensitive. The valid identifier *MiXeD* is the same as *mIxEd*.

Keywords are syntactic markers that provide a framework for event descriptions and serve to delimit user supplied parts of definitions. Keywords are case sensitive as well as reserved to the model builder. No keyword in any combination of character case may be used as an identifier. The following are the keywords of the *EDL* model building language.

cond **end** **event** **false**
is **primitive** **true** **with**

Type names are not keywords but identifiers that occur in limited contexts following event parameter and attribute name declarations. Currently valid type names are

integer - denoting positive or negative (or zero) whole numbers,

double - floating point numbers whose container size is the determined by the host system,

bitstring - whole unsigned integers,

time - values that denote time of day, used for event time stamps. Clock resolution is determined by the host systems, but the external representation is as described below,

string - contiguous lists of unsigned character-sized values that may contain both printing and non-printing values,

location - treated exactly like strings, but is used to location stamp event tuples. Interpretation of location types is reserved for later to accomodate different hardware and software architectures. For example, the location string might contain the processor number of a multiprocessor or a thread identifier of a lightweight process.

2.2 Constant Values

Constant expression values may be either numbers or character strings that are used to represent absolute quantities. Constant values are used in arithmetic and relational expressions and as indexes for event expression member events.

2.2.1 Numbers

EDL allows numbers to be specified as signed decimal and floating point numbers or radix-specific bit strings. The range of values, the container size, alignment, and representation are, as always, system dependent. Some examples

37.661	-392	0xFF01
0b101101	12.2e-2	0o77376

The keywords **true** and **false**, intended for use in relational expressions, represent the values 1 and 0, respectively.

Numbers consisting of a string of decimal digits (0..9) are typed as integers. Likewise, radix-specific bitstrings are classed as integers. Radix specified numbers begin with a zero followed by a radix specification character; then valid characters from the radix. Four common radixes are available:

binary - specified as **0bddd**; where each *d* may be 0 or 1,

octal - specified as **0oddd**; each *d* is 0 .. 7,

decimal - specified as **0oddd**; each *d* is 0 .. 9,

hexidecimal - specified as `0xddd`; each *d* is 0 .. 9 or (A ..f).

Other numbers (not radix-specific) may also begin with a zero character but are interpreted as decimal numbers. Floating point numbers can be written as decimal digit strings containing a decimal point (to separate the fractional part in the usual manner) or in scientific notation of a decimal string followed by a base ten exponent part. The decimal point or exponent part (or both) is required to distinguish floating point from integer numbers.

The general forms are:

```
{±}dd.dd
{±}dde{±}ddd
{±}dd.ddde{±}dd
```

Embedded spaces cause problems (e.g. between the *e* and the exponent value digit string) and the digits are, of course, all decimal digits.

2.2.2 Time values

Time values are an important part of events generated in complex systems. Event generators time stamp events with a form

```
[hours:]minutes:seconds[.hundredths]
```

EDL descriptions containing time typed values use this form to express constant time values that generally represent time differences or relative times. The *hours* and *hundredths* part of the time value are optional. Fifteen seconds is expressed 0:15; eight hours is expressed 8:0:0 .

2.2.3 Strings

A character string is written as a sequence of characters enclosed in quotation marks, ' " '. Non-printing characters can be inserted into string constants by writing the integer values preceded by a "\ " escape character. Some common non-printing characters have shorthand forms: \n, inserts a new-line character; \t, inserts a horizontal tabbing character; \e, inserts the ESC character; \b, inserts the backspace character (ctrl-h); \r, inserts a carriage return; \f, inserts a form feed character; \\, inserts the \ itself; \", inserts a double quote character (the normal string terminator). As with numbers, all specifics of the representation are system dependent. Some string examples

```
"vax9"
"An escape sequence \27[H"
```

Single character values common in most programming languages are supported only as integers. Single character values are expressed by sandwiching the single character between two ' marks, e.g., 'Q' or '\033'. Since integer and string types are incompatible types tool users should choose either a string or integer representation for single character values.

3. Directives to the Model builder

Inputs to the model builder can be supplied in two ways. The first is in a batch mode where a user simply supplies a file of *EDL* models and a library to be edited by the models. In this case, directives permit event files to be self contained units that specify libraries and other needed definition files. The other is a more interactive mode where the compiler reads from the standard input stream. For this case, directives control interaction with the associated event libraries, to flush the local definition cache or change the stream to other files.

Directives are instructions to the model builder to change its behavior in some manner. All directives to the model builder begin with a # character. The directive follows with no intervening space characters; arguments to the directive follow the directive string. All characters up to the end of line that are not used by the directive, are ignored.

3.1 File Inclusion

The model builder can be requested to stop reading the file it is currently in and begin reading an alternative file by using the `#include` directive. The include directive takes a single argument, a filename to switch attention to. The form of the include directive is

```
#include "filename.edl"
```

The model builder will read text from the included file until it encounters the end of file. Following end of file on the included file the model builder resumes reading from the file that contains the `#include` directive. Include directives can be nested to a maximum depth of eight files (one original, seven includes without encountering end of file).

3.2 Library Switching

The `#library` directive causes the model builder to close the currently open library and switch to a new library. All definitions in the cache are flushed to their library and the library information file is updated. Then the new library is read. The form of the directive is similar to the include directive,


```
#library "library-name"
```

will update and close whatever library is currently open then open the named library. The new library is located in the usual manner by concatenating the library name argument to the library root, given by either the environment variable `ELIBROOT` or a subdirectory, `~/EventLib`, located in the user's home directory.

3.3 Creating Background Libraries—Merge

It is sometimes desirable to combine two (or more) event libraries into a single viewpoint. For example separate libraries might exist that describe the process control functions and the I/O subsystem of an operating system. The `#merge` directive will read all events from its argument library and add them to the library being created.

```
#merge "library-name"
```

Libraries are mergeable only if all event names that each defines are unique. A merge will fail and produce no output if non-unique event names are encountered.

4. Describing High-Level Event Models

A single high-level event definition describes a model of some aspect of a system's overall behavior. Each event definition is composed, from a heading which supplies the name for the event class being defined and up to three kinds of defining clause which specify the structure and intent of the event. The `is` clause defines the event expression that specifies what event classes constitute the cluster and indicates acceptable orderings for their instances. The `cond` clause details a set of relational expressions defined over the attributes of the event expression event classes. These relational expressions constrain the attributes of events that appear in the event expression to specific values or value ranges. The event heading together with a set of attribute binding expressions defined in the `with` clause of a definition describe the event tuple for the event class represented by the event definition. The attribute binding expressions are also expressed in terms of the event expression members' attributes.

4.1 Event Heading

The event heading of an event definition associates a name by which the event class is known and referred to by users of the modelling tools employing *EDL*. The event name that appears in the heading is not necessarily the identifying name used internally by the *EBBA* model matching

components. However, it is important that there is available a unique mapping of system defined event names onto the external names with which an event view is constructed by a user. This is necessary so that users have some basis to believe that what they see represents their view of a system [Snodgrass85]. Within a given view, all event names are required to be unique.

An event heading may have an optional parameter list which provides a means of parameterizing events. The parameter list is specified as a list of names acting as place holders for values that are supplied when a request is made to match the event description to actual system activity. This permits a limited degree of tailoring an event recognition request to dynamic system conditions. The event heading of figure 1 introduces an event class named *FrontEndCompletion* which is invoked

```
event FrontEndCompletion( node:location ) is
```

Figure 1: Event heading with one parameter

with a single parameter, *node*. The *node* parameter can appear later as an operand in relational or arithmetic expressions in the **cond** or **with** clauses of the event description. A value is bound to this parameter when a request is made for recognition of an instance of the *FrontEndCompletion* event. Types will be described later in section 4.3. The example heading might be useful where a user is interested in instances of the *FrontEndCompletion* event only if they occur on a specific node of the distributed system. For example, in the debugging monitor, a request for recognition of an instance of the *FrontEndCompletion* would have the form⁴:

```
Match-> (recognize (FrontEndCompletion "vax9"))
```

The event recognizer would be requested to watch the event stream for an instance of the high-level *FrontEndCompletion* event with the string value "vax9" bound to the *node* parameter.

4.2 Event Expression

The event expression is a regular expression-like string of event symbols and event operators. The event symbols represent constituent behaviors of the high-level event model which is described by the enclosing event description. The operators indicate alternative orderings that sets of event instances may have in order to match the event expression members. In general, the event expression will only describe an outline for a behavior model. The event operators define acceptable orderings, thus implicitly imposing time relations on event expression constituent events. However, events

⁴The "Match->" is a prompt from the Event Monitor toolset component indicating that it is looking for something to do.

carry a richer set of attributes that may be used to relate collections of events. These relations are handled by the `cond` clause, described later.

Figure 2 expands the earlier *FrontEndCompletion* example to include an event expression. The *FrontEndCompletion* event is defined in terms of three other events: *InstantiateKsi*, *CreateHyp*, and *InvokeKs*.

```

event FrontEndCompletion(node:location) is
    InstantiateKsi[1] ◦ CreateHyp* ◦ InstantiateKsi[2] ◦ InvokeKs
end

```

Figure 2: Event heading with event expression

The event expression provides a means to perform coarse filtering of event stream instances. When an attempt is made to recognize an instance of the event, only event stream instances in the same class as event expression symbols need be examined for possible inclusion as constituent events.

The subscript notation appended to each of the *InstantiateKsi* event symbols serves to distinguish the two occurrences when they are used as operand qualifiers in `cond` and `with` clauses. Without the subscripts all expression references to attributes of an *InstantiateKsi* event would access the values contained in the instance bound to the first *InstantiateKsi* of the event expression.

Event expressions need to describe a wide range of behavioral patterns. In distributed systems, behaviors will consist of sets of events which are only partially ordered. Sometimes, the set consists of events which need to be ordered in strict sequences. Other times the events occur concurrently, and will be observed in arbitrary order. Conventional regular expression formalisms are capable of describing sequences in various ways but do not provide means for expressing concurrency. *EDL* event operators include operators that specify sequencing, choice and repetition as well as an operator describing concurrently occurring sets of events.

4.2.1 Sequencing

The binary sequence operator, indicated with “◦” (@ on some keyboards), specifies that its operand events are to occur in the left to right order in which they are written. For example the partial expression

```
... InstantiateKsi ◦ InvokeKs ◦ SendHyp ...
```

indicates that the described behavior consists of an instance of the *InstantiateKsi* event followed at a later time by an *InvokeKs* event, which is subsequently followed by a *SendHyp* event. Implicitly, the three events are related by a relation of their time attributes where

$$\textit{InstantiateKsi.time} < \textit{InvokeKs.time} < \textit{SendHyp.time} .$$

4.2.2 Repetition

Indefinitely long sequences of events are specified by two postfix unary operators, plus and star, indicated by $+$ and $*$, respectively. Plus indicates that the described event sequence consists of one or more repetitions of its event operand. Star indicates a sequence of zero or more instances of its operand event. Thus, the partial event expression

$$\dots \textit{InstantiateKsi} \circ \textit{CreateHyp}^* \circ \textit{InstantiateKsi} \dots$$

describes the sequence

```
(InstantiateKsi . . .)
(InstantiateKsi . . .)
```

or

```
(InstantiateKsi . . .)
(CreateHyp . . . )
(InstantiateKsi . . .)
```

or

```
(InstantiateKsi . . .)
(CreateHyp . . . )
(CreateHyp . . . )
(CreateHyp . . . )
(CreateHyp . . . )
(InstantiateKsi . . .)
```

and so on. Whereas

$$\dots \textit{InstantiateKsi} \circ \textit{CreateHyp}^+ \circ \textit{InstantiateKsi} \dots$$

will describe all but the first of the above event sequences.

4.2.3 Choice

Specifying a choice between different events is accomplished by the alternation operator, “|”. Alternation indicates that an instance of at least one of its operand events is necessary to match the event expression. For example, the partial event expression

$$\dots \text{CreateDDGoal} \circ (\text{ReceiveHyp} \mid \text{CreateHyp}) \circ \text{InvokeKs} \dots$$

matches the event subsequence

```
(CreateDDGoal . . .)
(CreateHyp . . .)
(InvokeKs . . .)
```

or the subsequence

```
(CreateDDGoal . . .)
(ReceiveHyp . . .)
(InvokeKs . . .)
```

Random choice among a group of operand events is possible by specifying a form:

$$\text{CreateHyp} \mid \text{SendHyp} \mid \text{CreateGoal} \mid \dots \mid \text{SendGoal}$$

The above expression indicates that the specified behavior consists of any *one* of the *CreateHyp...SendGoal* events. The alternation operator also has the property that the fully parenthesized expression

$$(\dots((\text{CreateHyp} \mid \text{SendHyp}) \mid \text{CreateGoal}) \mid \dots \mid \text{SendGoal})$$

is the same as the above, unparenthesized expression.

Note that the definition of alternation does not state “... but not both ...” – exclusion is not its purpose. The nature of the event stream is such that all alternates could occur in the proper context. However, to successfully match the event expression to the stream only requires one of the alternates. Which event is chosen will depend on the tools employed to match the models to event streams.

4.2.4 Concurrency

Concurrency is described using the shuffle operator, written “ Δ ” (\wedge on some keyboards). The shuffle operator designates a set of events, *all* of whose members must occur, but there is no preferred ordering imposed on the set members. For example,

$$\dots \text{InvokeKs} \circ (\text{CreateHyp} \Delta \text{SendHyp}) \circ \text{InvokeKs} \dots$$

describes the sequence

```
(InvokeKs . . . .)
(SendHyp . . . .)
(CreateHyp . . . .)
(InvokeKs . . . .)
```

or the sequence

```
(InvokeKs . . . .)
(CreateHyp . . . .)
(SendHyp . . . .)
(InvokeKs . . . .)
```

Similar to the alternation operator, a series of event names joined with shuffle operators in an event expression string such as

$$\dots \text{CreateHyp} \Delta \text{ReceiveHyp} \Delta \text{CreateDDGoal} \dots$$

is the same as the parenthesized version

$$\dots ((\text{CreateHyp} \Delta \text{ReceiveHyp}) \Delta \text{CreateDDGoal}) \dots$$

This use of the shuffle operator designates a cluster of events in which *all* of the events participating in the shuffle must occur.

The use of the shuffle operator to denote concurrent events does not imply that (in the above example)

$$\text{CreateHyp.time} = \text{ReceiveHyp.time} = \text{CreateDDGoal.time}$$

It is also not true that if one operand of a shuffle occurs “first”, the rest will occur “later”. That is, if the following are true about the constituents bound to the event instance represented by the above shuffle

$$\begin{aligned} \text{CreateHyp.time} &= t_1 \\ \text{ReceveHyp.time} &= t_2 \\ \text{CreateDDGoal.time} &= t_3 \end{aligned}$$

there is no guarantee that an event sequence matching the shuffle event string implies any relations between the attribute times t_1 , t_2 , or t_3 .

```

event FrontEndCompletion( atnode:location ) is
    InstantiateKsi[1] ◦ CreateHyp* ◦ InstantiateKsi[2] ◦ InvokeKs
cond
    InvokeKs.location == atnode;
    InstantiateKsi[1].location == InstantiateKsi[2].location
end

```

Figure 3: Event definition with constraining clauses

4.3 Cond Clauses

The **cond** clause contains a set of relational expressions which are used to constrain the attributes of event expression members to certain values or value ranges. The set of **cond** clause relationals that are part of an event description specify that only instances of event expression members possessing certain characteristics are to be considered eligible to match the described high-level behavior.

The relations among events expressed by **cond** clause constraints provide finer grained filtering than that provided by the event expression since they are defined in terms of the attributes of event expression members. For recognition of a particular event, the high-level event recognition machinery only selects, from the event stream, event instances that match event expression members. Following this coarse filtering, any **cond** clause relations defined in the high-level event involving attributes of the instance matching an event expression member must be evaluated to determine if the selected instance possesses appropriate attributes.

These constraints on event attributes thus serve both to narrow the scope of a model and as filters which eliminate unnecessary instances of events from inclusion into a behavior model. A **cond** clause is an optional part of an event description. In the absence of a **cond** clause, any set of events from the appropriate classes which match a pattern prescribed by the event expression will constitute an instance of the defined event. Figure 3 is the earlier *FrontEndCompletion* example extended to include **cond** clause constraints.

The set of operators available for **cond** clause relational expressions is a rich set of arithmetic, logical, and relational operators. Arithmetic operators include addition (+), subtraction (-), multiplication (*), division (/), and modulus (%). Bitwise logical operations are shift left (<<), shift right (>>), and (^^), and or (||). The arithmetic and bitwise logicals are all binary numeric operations and return numeric values. Available relational operators are the usual less than (<), less than or equal (<=), equality (==), inequality (!=), greater than or equal (>=), and greater

than ($>$). These are defined over all values and return boolean **false** (value 0) and **true** (value 1) results. Logical connectives and ($\&$) and or (\mid) supply conjunction and disjunction of boolean values. Unary complement operators are negate ($-$) for arithmetic values, bitwise logical complement (\sim), and not ($!$) for boolean expressions. Of course, all of these may be augmented with external functions defined in the *EBBA* tool environment.

Expression operands have five sources: event heading parameters, event instance attributes, so-called “lexical” functions, system defined functions and simple constant values. Explicit means is provided to associate data type information with all operands. Thus, type compatibility checking is performed for individual expressions to ensure they can actually be evaluated meaningfully.

No capability is provided for creating data types. Users are limited to several scalar types generally considered useful

numbers - can take on type *integer*, *double* (floating point), and *bitstring*

strings - have a single data type, *string*.

time - a specific time type is provided, *time*. The written value for time has the form

[hours:]minutes:seconds[.hundredths]

The hours and hundredths parts are not necessary for a complete time specification.

location - provision exists for values to be typed as *location* although currently location and string types are considered the same. The general convention for designating locations has values take the form

"node-name: component-name"

The following subsections detail each source of expression operand by describing the access mechanism and when an operand takes on a value. This discussion of operands and value-producing expressions applies equally well to attribute binding expressions in **with** clauses.

4.3.1 Event Heading Parameters

Event heading parameters are identifiers listed in the parameter section of an event heading. Their importance is that they can be used to parameterize the value producing expressions in an event definition. Each parameter in the list is declared with its type using the

parameter-name: type

notation. Event heading parameters are accessed in expressions as unadorned identifiers. Values are bound to event heading parameters when a recognition request is made for an instance of the event class. In **cond** clause relational expressions, use of event heading parameters serves to tailor

the filtering effect of the `cond` clause to dynamic system conditions. Refer to figure 3 for an example of the use of an event parameter as an expression operand.

4.3.2 Event Attributes

Event attributes are the most important kind of valued operand because event attributes are the primary means for focusing particular behavior abstractions. Values represented by these operands are obtained from event instance tuples bound to event expression members. Event attribute operands are specified by using the attribute name prefixed by the name of the event with which the attribute is associated. The qualifying event symbol must appear as a member of the event expression of the enclosing event description.

If the *FrontEndCompletion* event expression members have the event templates

(InstantiateKsi ...time location)

(CreateHyp time ...location)

(InvokeKs time ...location)

the focus of the *FrontEndCompletion* event can be narrowed by specifying constraints in terms of the event expression member attributes. The attribute access is written using the common field access dot notation as in numerous programming languages [Horowitz84]. For example, the event definition of figure 3 constrains all of the event stream events which are bound to the event expression members to have occurred on the same node of the system.

Recall that the subscript notation is important when an event name is used more than one time in an event expression. Value expressions use the subscripts to specify which qualifying event supplies the attribute value when the expression is evaluated. The subscript notation only applies to the explicit mentions of an event in an event expression. The correspondence between event expression events and qualified event attributes is made when the event description is created. The implication of this is that the attributes of a repetitive event instance, e.g. the *CreateHyp* events bound to the event expression

$\dots \circ \textit{CreateHyp}^+ \circ \dots$

are **not** accessed as *CreateHyp[1].node*, *CreateHyp[2].node*, etc. The only valid subscripted event attribute qualifier for this string would be *CreateHyp[1].node*, since only a single mention of the *CreateHyp* event is made in the event expression. Accesses to specific instances of the string of *CreateHyp* events bound to the *CreateHyp⁺* event expression member are performed through lexical functions.

```

event FrontEndCompletion( count:integer ) is
  InstantiateKsi[1] ◦ CreateHyp* ◦ InstantiateKsi[2] ◦ InvokeKs
cond
  numberOf(CreateHyp) < count
end

```

Figure 4: Event expression “Lexical” functions

4.3.3 Event Expression “lexical” Functions

Certain event expression constructs have attributes which result from the binding of event stream instances to the event expression member events. Attributes include the total number of events bound to a repetitive event expression or the absence of an event instance from the bound event string. The first is given by the *numberOf()* lexical function, the later by the *undefined()* lexical function. For example, as shown in figure 4 the *FrontEndCompletion* event might only be “interesting” if the number of *CreateHyp* events is smaller than some number, *count*, supplied as a parameter. Other types of lexical functions help to unravel things which are not easily expressed in the event expression format, such as accessing the first or last instance of a repetitive sequence. For example the function

$$\textit{last}(\textit{CreateHyp}).\textit{node}$$

selects the *node* attribute of the last *CreateHyp* instance which has been bound to the *CreateHyp*⁺ part of the event expression string.

4.3.4 Environment Specific Functions

These functions are supplied to return values obtained from the environment supporting the abstraction of event instances. What functions are supplied is highly implementation dependent, but they at least include a *timestamp()* and a *locationstamp()* function. Environment specific functions are provided largely as a catch-all for unanticipated, but necessary, functionality in the evaluation environment. Functions are defined using the *EBBA* toolset Extension Language (*elll*) and are loaded into the Event Recognizer when it starts either through startup files or interactively as the recognizer is executing.

4.4 With Clauses

The list of names for attributes possessed by an event and a description of how to bind values to each attribute when an instance of the event is recognized is found in the **with** clause expression. Each **with** clause expression names and describes a single event attribute. The attribute is defined as an identifier together with the attribute's type which is the target of an assignment operator. The expressions which provide values for event attributes are composed from the same set of operators and operands as **cond** clause expressions.

All events have predefined attributes *time* and *location*. By convention, these attributes fill in the last two attribute slots of an event tuple. This can be overridden by explicit inclusion of the *time* attribute as

```
time:time := timestamp()
```

and the *location* attribute as

```
location:location := locationstamp()
```

in of the desired position in the **with** clause.

When a set of events matching one of the possible event expression sequences has been accumulated and all **cond** clause constraints evaluate satisfactorily, the event is instantiated in its own right. Using the context supplied by its parameters and attributes of the constituent events, the attribute binding expressions are evaluated and the value for each slot in the event tuple template is filled in.

To illustrate, in figure 5 the familiar *FrontEndCompletion* event can be expanded to possess

```
event FrontEndCompletion( count:integer ) is
    ...
with
    hypcount:integer := numberof(CreateHyp);
    elapsedtime:time := InvokeKs.time - InstantiateKsi[1].time
end
```

Figure 5: Example attribute binding expressions

attributes (in addition to the predeclared attributes for *time* and *location*) for *elapsedtime* and *hypcount*. The *elapsedtime* attribute is bound to the difference in the time attributes of the events which are first and last in the event string bound to an instance of the *FrontEndCompletion* event.

The *hypcount* attribute reflects the length of the string of *CreateHyp* instances bound to this event expression.

Like the **cond** clause, the **with** clause is an optional part of the event definition. However, every event instance will carry with it certain predefined attributes, such as time of occurrence, that might serve to distinguish various instances of the class. Other predefined attributes might be dependent on specific characteristics of the system, such as the name of the processor node on which the event occurred.

5. Describing Primitive Events

Primitive events, those events that represent the lowest level of observable behavior, are also defined by *EDL* descriptions. The description for a primitive event is an abbreviated one which contains no event expression, **cond** clause, or parameter list. Each of these parts of an event description is related to some active property of event recognition. Primitive events are created by the system which is under investigation and are characteristic of that system. They are atomic entities that a user of model specification and matching tools has no control over. If a user desires to see only primitive event instances with certain attributes, it is necessary to construct a high-level event that expresses the appropriately constrained cases.

Primitive event descriptions consist of a special event heading and a reduced form of **with** clause. The event heading names an event class and declares the event class as **primitive**. Following the **primitive** keyword is a value, the system specific identifier, that the event generating source puts into the event class slot of an instance of this event class. The form and value of the system specific identifier following the **primitive** keyword is system and implementation dependent. It could be a number, a string, or some arbitrary bit pattern.

The unique system identifier has an effect of permitting flexible binding of names to actual events. It is sometimes necessary to shift the viewpoint on a system to a "lower" level of primitive event. The capability in *EDL* to easily rebind primitive events facilitates this type of rebinding of event to event generator.

In a primitive event definition, the set of attributes defined by the primitive event must be listed as for high-level events.

Also, the time and location attributes are predefined and occupy the last two event tuple attribute slots. However, there are no code strings associated with the attributes. Event attributes result from action by the event generator that is, in the primitive event case, an unaccessible system activity. The *EDL* primitive event definition thus defines the tuple for a characteristic primitive event and provides a handle to match it with system generated events. Figure 6 shows the relationship between an *EDL* primitive description, the template for the event tuple, and a few

```

event CreateHyp is primitive "CREATE_HYP"
with
    hypname
end

(CreateHyp, hypname, time, location)

(CREATE_HYP "h:09:0034" 57391 "VAX9")
(CREATE_HYP "h:02:0173" 159318 "VAX2")

```

Figure 6: Primitive description, event template, and instances

```

event FrontEndCompletion( atnode:location ) is
    InstantiateKsi[1] ◦ CreateHyp* ◦ InstantiateKsi[2] ◦ InvokeKs
cond
    InvokeKs.node == atnode;
    InstantiateKsi[1].location == InstantiateKsi[2].location
with
    hypcount:integer := numberof(CreateHyp);
    elapsedtime:time := InvokeKs.time - InstantiateKsi[1].time
end

```

Figure 7: Complete *FrontEndCompletion* event definition

instances of the primitive event. All that is necessary to use this primitive event is that the system generate some event tuple which matches the event definition.

6. A Complete Example

Now that all of the parts of a high-level event description have been explained it might be useful to illustrate many of the views of an event definition. A complete *FrontEndCompletion* event might be described as in figure 7. This description represents the event tuple

(*FrontEndCompletion hypcount elapsedtime time location*) .

The implicit attributes *time* and *node* are always present and are part of every event tuple. Some example instances of this event could be

```
(FrontEndCompletion 25 371 233:83 "vax9")
```

and

```
(FrontEndCompletion 4 54 184 "vax6")
```

If the observed event stream contained a sequence of events

```
(InstantiateKsi . . . 10:0 "vax7")
(InstantiateKsi . . . 9:0 "vax6")†
(InvokeKs . . . 12:0 "vax7")
(CreateHyp . . . 13:0 "vax7")
(InvokeKs . . . 121:0 "vax6")
(CreateHyp . . . 146:0 "vax6")†
(CreateHyp . . . 172:0 "vax7")
(InstantiateKsi . . . 223:0 "vax6")†
(InvokeKs . . . 250:0 "vax6")†
```

then an instance of the *FrontEndCompletion* event derived from this stream would have the events marked with † as constituents. The instantiated *FrontEndCompletion* event could be

```
(FrontEndCompletion 1 241:0 250:0 "vax6")
```

The *FrontEndCompletion* event can also be viewed in terms of its structure, which is derived from its constituent events, their constituents, etc. Since self-referent event descriptors are not allowed, the structure defines a tree with the given high-level event as the root. If the *FrontEndCompletion* event is incorporated into a higher level event, *FirstSigActivity* with the definition of figure 8 the

```
event FirstSigActivity( atnode:location ) is
  FrontEndCompletion o (ReceiveHyp | CreateHyp)
cond
  CreateHyp.node == atnode;
  ReceiveHyp.node == atnode;
  FrontEndCompletion.node == atnode
end
```

Figure 8: A higher level event

structural representation of *FirstSigActivity* would be as shown in figure 9.

This structural view is useful when trying to understand the overall structure of an event based behavior description. With proper annotation this representation might make it easier to judge

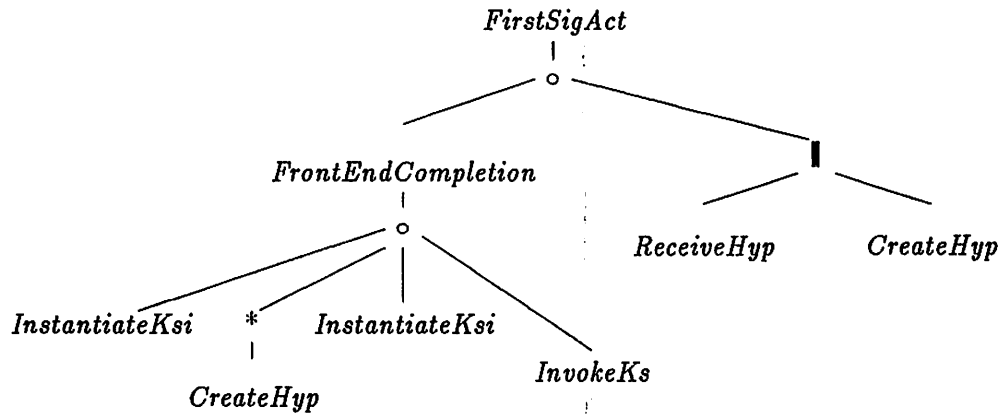


Figure 9: Structural representation of a high-level event

how well this abstraction matches the actual system activity. Also, the structural view aids in analyzing certain difficulties with the abstraction process itself.

A. Using the Model builder

After it is determined where the executable form of the model builder resides it is handy to create an alias (`ed1`) to be used to invoke it.

```
ed1 [-o] [-t] [-l libname] [-m libname] [-e event] [file1 file2 ...]
```

file1, *file2*, ... are text files containing *EDL* definitions and modeler directives that are to be translated into recognizers. If no files *file1*, *file2*, ... are specified, then the model builder has been invoked in interactive mode and will attempt to read from standard input, which can be attached to almost anything. The indicated files are read in the order they are given on the invocation line and the events they contained are processed in the order they appear in the source files. These files may contain any of the model builder directives described in Section 2 to change the behavior of the model builder. Be careful.

The `-l` directive specifies an initial event library (or librarian) that maintains the viewpoint that the compiled events are related to. The *libname* parameter names a previously created event definition library (or one about to be created). The model builder will access the library indicated by the `-l` option, then use the definitions contained in *file1*, *file2* ... to edit the opened library. If no library is specified (no `-l` option selected) the library named "default" is used. A library may be specified or changed in a `#library` modeller directive if not specified at invocation time.

The `-m` option specifies a library that is to be merged with the main library. Any number of merge libraries may be specified with the proviso that the events in the main and merged libraries do not conflict. Each merged library event will be checked for uniqueness and if no problems arise will be added to the main library. More than one merge library may be specified with multiply `-m` options.

If `-o` is specified, no library outputs will be created by the compiler. Syntax and consistency checking is performed but no new or changed events are written back to the event library. The base library for the compiled events must be available to allow examination of high-level event constituents.

The `-e event` option indicates that only the named *event* is to be translated. All events in the file list should be syntactically correct, since selection for translation is performed after parsing of individual event models (and the selected event might get caught up in the dumb error recovery technique).

Normally *time* and *location* fields are automatically added to any event tuple that does not explicitly include them. The `-t` option tells the model builder to NOT supply default *time* and *location* fields for events that do not have them specified.

B. Model Builder outputs

The model builder prints some diagnostic and informative messages as it executes. Less exciting outputs are the translated event descriptions that may be used by the event recognizer to recognize model instances. Also output by the model builder are files that aid primitive event generation.

B.1 Diagnostic and informative messages

The following list of messages is output by the Compiler and written to the standard error stream. Most of them are self-explanatory, and will require that the source text of an event be edited and compiled again. Error recovery is *very* dumb, so innocuous syntax errors can be quite amusing.

Attempt to redefine formal parameter *param*

The named event heading parameter was defined more than one time.

Attempt to redefine primitive event *event-name*

Primitive events may not be redefined. The user should clear the library using the delete files shell command or request the librarian interface for the library delete the the redefined primitive event.

Duplicate attribute name *attribute-name*

An attempt was made to redefine a primitive event attribute.

Attempt to redefine attribute *attribute-name*

An attempt was made to redefine the named high-level event attribute.

No event number for event *event-name*

A request was made to the librarian for an event number for the named event, but none was returned. Either more than 2^{16} high-level events have been defined or there is an internal error.

Unknown Event *event-name*

An event name which has not been defined yet appears as an operand in an event expression.

Too many/Insufficient parameters to constituent *event*

An event expression member is a high-level event with parameters, and the number of actual parameter bindings is incorrect.

Undeclared local variable *name* in expression

A simple name that is not declared as an attribute or formal parameter is being used in a **cond** or **with** clause expression.

Event *event* not in event expression.

An event name that is not used in the event expression for the event is used to qualify an attribute used as an operand in a **cond** or **with** clause expression.

Attribute *name* not part of event *event*.

The attribute *name*, qualified by the *event* is not defined by the *event*.

Invalid numeric text to decode '*text*'

This is probably an internal error, it indicates that what was once considered to be valid numeric text cannot be translated into a binary form. The text is probably the index of a qualified event.

Character *char* does not begin a token.

The indicated character is not a valid prefix for a symbol in the EDL language. This error will be followed by zero or more messages regarding recovery attempts.

Include files nested too deeply, file *file-name* ignored.

Only 8 levels of source file can be maintained by the compiler (one original plus 7 #include'd files). The indicated file tried to be the 9th.

Newline before terminator (*char*) for string

char was used to begin a character string, but no matching character was found before the end of the source text line.

Cannot open include file *file-name*.

The named file could not be opened for reading, either it does not exist or one of the directories in its file path is inaccessible.

Unknown option '*string*' specified.

The command that invoked the compiler contained an unrecognized option, or a previous option did not have a value, but was given one.

Start state mapping is missing a routine start.

This is an internal error. The Shuffle Automata contains a number of sub-machines, one of which has not had its starting state added to the list of machine start states.

Syntax error near *token-name*.

Dumb error recovery in action. There was a syntax error at the token indicated. This error will be followed by zero or more messages regarding recovery attempts.

Unexpected End-of-File encountered.

The end of the source text was found before the current event definition was completed.

- Token *token-name* deleted from input.

More dumb error recovery. When a syntax error is encountered, symbols are deleted from the input stream until something the compiler likes better comes along.

Invalid ASTnode op *number*.

This is an internal error. The abstract syntax tree constructed to represent an event has an unrecognized value in its node type field.

Attempt to insert duplicate symboltable entry for *name*.

This is an internal error. A routine attempted to add a name to the local symbol table for an event without checking to see if it was already there.

Unknown node type in ConstEvalExpr.

This is an internal error. An arithmetic or relational expression has an invalid node type in its abstract syntax tree.

B.2 Real Compiler Outputs

The compiler reads event source text and events stored in libraries and outputs libraries and a file of primitive event definitions. As a matter of organization it is probably reasonable to put all of the source for the primitive events for a viewpoint (or library) into a single file. The primitive events for a viewpoint cannot be redefined, so when they are syntactically correct they are more or less a run once and kept around for reference. Each high-level event should get its own file container, they are changeable and the librarian/model-builder interface can be cleaner – but this is not a requirement.

The files created as a result of a compilation of a collection of primitive and high-level event descriptions include the following:

- a library directory that is located in the library root directory given by either the environment variable `ELIBROOT` or a subdirectory, `~/EventLib`, located in the user's home directory. This file will have the name `~/EventLib/lib`.
- a library information file `lib.INFO` that holds the library status (high-level and primitive event-number servers, counts, etc....) The name for this file will be `~/EventLib/lib/lib.INFO`.
- a primitive event formatting file `lib.h` that contains descriptors for the primitive events. This file is used by the default primitive event generator in `libBA` to put together event tuples. The final name for this file is `~/EventLib/lib/lib.h`
- as many files of the form `nnnnn.EVT` as are needed for the translated event descriptors. `nnnnn` is $1 \dots 2^{16} - 1$ for primitive events; and $2^{16} \dots 2^{\text{MAX_EventNumber}}$ for high-level events. (`MAX_EventNumber` can be found in the `debug/include/params.h` file if the reader is really interested. These files will have the final name `~/EventLib/lib/nnnnn.EVT` There may be more high-level event files than there are high-level events since each time a high-level event is compiled a new event number is issued for it. The librarian utility (`1u` or the librarian/model-builder interface) or the `delete files` command can be used to clean up those unwanted event files.

The details of the file contents are found in another report

P.C. Bates, "The EBBA Toolset Implementation Notes"

Also, detailed descriptions of the compiler internals and operation are found in that report.

C. BNF Description of EDL

event_defs	::=	event_description event_defs event_description
event_description	::=	event event_heading is_clause cond_clause with_clause end
event_heading	::=	identifier identifier (arglist)
arglist	::=	formal_id arglist , formal_id
formal_id	::=	identifier:type
is_clause	::=	is event_expression
event_expression	::=	re_expr primitive value
re_expr	::=	re_sexpr re_expr re_sexpr
re_sexpr	::=	re_term re_sexpr \circ re_term
re_term	::=	re_factor re_term Δ re_factor
re_factor	::=	constituent_event re_factor repetition (re_expr)
constituent_event	::=	identifier elist event_index
elist	::=	(expr_list) ϵ
event_index	::=	[number] [identifier] ϵ
repetition	::=	* +
with_clause	::=	with attribute_list ϵ
attribute_list	::=	attribute attribute_list ; attribute
attribute	::=	attribute_name attribute_name := expression

attribute_name	::=	identifier:type
expression	::=	primary - expression ! expression ~ expression expression binop expression
primary	::=	value qualified_name identifier (expression) identifier (expr_list)
expr_list	::=	expression expr_list , expression
value	::=	number string boolean
boolean	::=	true false
binop	::=	* / % + - << >> < <= >= > == != & &&
qualified_name	::=	identifier event_index . identifier
cond_clause	::=	cond boolean_exprlist ϵ
boolean_exprlist	::=	expression boolean_exprlist ; expression
string	::=	" characters "
characters	::=	character characters character
character	::=	ASCII-character \number
number	::=	bin_num oct_num decimal hex_num char_num
bin_num	::=	0b bin_string
bin_string	::=	bin_digit

		bin_digit bin_string
bin_digit	::=	0 1
oct_num	::=	0o oct_string
oct_string	::=	oct_digit oct_digit oct_string
oct_digit	::=	0 1 2 3 4 5 6 7
decimal	::=	dec_string 0d dec_string
dec_string	::=	dec_digit dec_digit dec_string
dec_digit	::=	0 1 2 3 4 5 6 7 8 9
hex_num	::=	0x hex_string
hex_string	::=	hex_digit hex_digit hex_string
hex_digit	::=	0 1 2 3 4 5 6 7 8 9 A B C D E F
type	::=	integer double time string location
char_num	::=	'character'