

FLAME
Flexible Large-scale Agent-based Modelling
Environment
User Manual

Mariam Kiran

University of Sheffield

Date updated

April 2010

Abstract

FLAME (Flexible Large-scale Agent-based Modelling Environment) is a tool which allows modellers from all disciplines, economics, biology or social sciences to easily write their own agent-based models. The environment is a first of its kind which allows simulations of large concentrations of agents to be run on parallel computers without any hindrance to the modellers themselves. FLAME was employed in the EURACE project for writing various economic agent-based models.

This document present a comprehensive guide to the keywords and functions available in the FLAME environment for the modellers to write their own agent models to facilitate their research. This user manual describes how to create a model description and write implementation code for the agents.

Installation guides to the tools required with FLAME have been provided separately in the document titled '*Getting started with FLAME*'. This documentation also contains details on executing example FLAME code and running your own models. This user manual is written in accordance with Xparser version 0.16.2.

FLAME has been produced and is owned by the University of Sheffield under copyright laws. It is used in a variety of projects for agent-based modelling like EPITHELIUM (tissue modelling), EURACE (economic modelling) and SOCIAL CAPITAL (social science) modelling.

Contents

1	Introduction	6
2	Working with FLAME	7
2.1	X-Machine	8
2.2	Swarm Example	9
2.3	Transition Functions	10
2.4	Memory and States	11
3	Model Description	12
3.1	Tags within the XModel	13
3.2	Model in Multiple XML Files	14
4	Environment	15
4.1	Constant Variables	15
4.2	Source C Files	16
4.3	Time Units	16
4.4	Data Types	17
5	Agents	18
5.1	Agent Memory	19
5.2	Agent Functions	19
5.3	Birth and Death of Agents	20
5.4	Function Conditions	21
5.4.1	Comparison Operators	21
5.5	Time conditions	22
5.6	Messages in and out of Functions	23
5.7	Message Filters	23
6	Messages	25
7	Model Implementation	25

7.1	Accessing Agent Memory Variables	26
7.1.1	Using Model Data Types	26
7.1.2	Using Dynamic Arrays	27
7.2	Sending and receiving messages	28
8	Model Execution	28
8.1	Xparser Generated Files	28
8.2	Start States Files 0.xml	30
9	Installation Notes on FLAME	31
9.1	MinGW	32
9.2	Libmboard	32
9.3	Extra Useful Installations	33
9.3.1	GDB	33
9.3.2	Dotty	33
10	Running a Simulation	33
10.1	Running Job Scripts	34
11	Notes for the Programmers of FLAME	36
11.1	Memory Allocation and its Problems	36
11.1.1	Dynamic Arrays	36
11.1.2	Data Types	36
11.1.3	Agent Memory Management	36
11.2	Agent Execution	37
11.3	Communication	38
11.4	Libmboard	39
12	XParser Distribution	39
13	XParser generated files for FLAME programmers	39
14	XParser Versions	43

15 Example Models

43

A XML DTD

43

1 Introduction

FLAME (Flexible Large-scale Agent-based Modelling Environment) is a tool which allows modellers from all disciplines, economics, biology or social sciences to easily write their own agent-based models. The environment is a first of its kind which allows simulations of large concentrations of agents to be run on parallel computers without any hindrance to the modellers themselves.

The FLAME framework is a tool which enables creation of agent-based models that can be run on high performance computers (HPCs). The framework is based on the logical communicating extended finite state machine theory (X-machine) which gives the agents more power to enable writing of complex models for large complex systems.

The agents are modelled as communicating X-machines allowing them to communicate through messages being sent to each other as per designed by the modeller. This information is automatically read by the FLAME framework and generates a simulation program which enables these models to be parallelised efficiently over parallel computers.

The simulation program for FLAME is called the **Xparser**. The Xparser is a series of compilation files which can be compiled with the modeller's files to produce a simulation package for running the simulations. Various tools have to be installed with the Xparser to allow the simulation program to be produced. These have been explained in the accompanying document, '*Getting started with FLAME*'.

Various parallel platforms like SCARF, HAPU and Iceberg, have been used in the development process to test the efficiency of the FLAME framework. This work was done in conjunction with Science and Technology Facilities Council at Rutherford Appleton Laboratories¹.

A block diagram of the FLAME files has been presented in Figure 1 where,

Model.xml contains the whole structure of your model which is the agent descriptions, memory variables, functions, messages.

Functions.c contains the implementations of the functions specified in the .xml file.

0.xml contains the initial states of the memory variables of the agents which is the initialisation of all parameters.

The number of the resulting XML files depends on the number of iterations you specify to run your model (through Main.exe).

This document is a comprehensive guide to modellers of all disciplines to write their own agent models with ease. The key advantage of FLAME lies in the fact that modellers from

¹More details of the results obtained can be found in '*Deliverable 1.4: Porting of agent models to parallel computers*'.

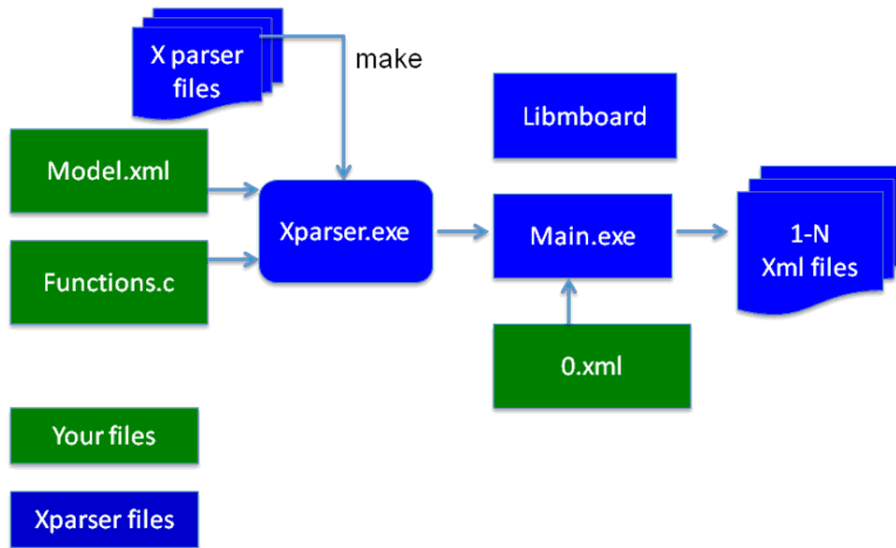


Figure 1: Block diagram of the Xparser, the FLAME simulation program. Blocks in blue are the files automatically generated. The green blocks are modeller files.

all disciplines, regardless of belonging to the computer science background can write their own models easily and run parallel simulations of them.

2 Working with FLAME

Traditionally software behaviour has been specified using finite state machines to express its working. Extended finite state machines (X-machines) are more powerful than the simple finite state machine as it gives the model more flexibility than a traditional finite state machine.

FLAME uses X-machines to represent all agents acting in the system. Each would thus possess the following characteristics:

- A finite set of internal states of the agent.
- Set of transitions functions that operate between the states.
- An internal memory set of the agent.
- A language for sending and receiving messages among agents.

2.1 X-Machine

An X-machine is defined as,

$$X = (\Sigma, \Gamma, Q, M, \Phi, F, q_0, m_0) \quad (1)$$

where,

- Σ are the set of input alphabets
- Γ are the set of output alphabets
- Q denotes the set of states
- M denotes the variables in the memory.
- Φ denotes the set of partial functions ϕ that map and input and memory variable to an output and a change on the memory variable. The set $\phi: \Sigma \times M \longrightarrow \Gamma \times M$
- F in the next state transition function, $F: Q \times \phi \longrightarrow Q$
- q_0 is the initial state and m_0 is the initial memory of the machine.

Figure 2 shows the structure of how two X-machines will communicate. The machines communicate through a common message board, to which they post and read from their messages. Using conventional state machines to describe the state-dependent behaviour of a system by outlining the inputs to the system, but this failed to include the effect of messages being read and the changes in the memory values of the machine. X-Machines are an extension to conventional state machines that include the manipulation of memory as part of the system behaviour, and thus are a suitable way to specify agents. Describing a system in FLAME includes the following stages:

- Identifying the agents and their functions.
- Identify the states which impose some order of function execution with in the agent.
- Identify the input messages and output messages of each function (including possible filters on inputs which will be explained in Section 5.7).
- Identify the memory as the set of variables that are accessed by functions (including possible conditions on variables for the functions to occur).

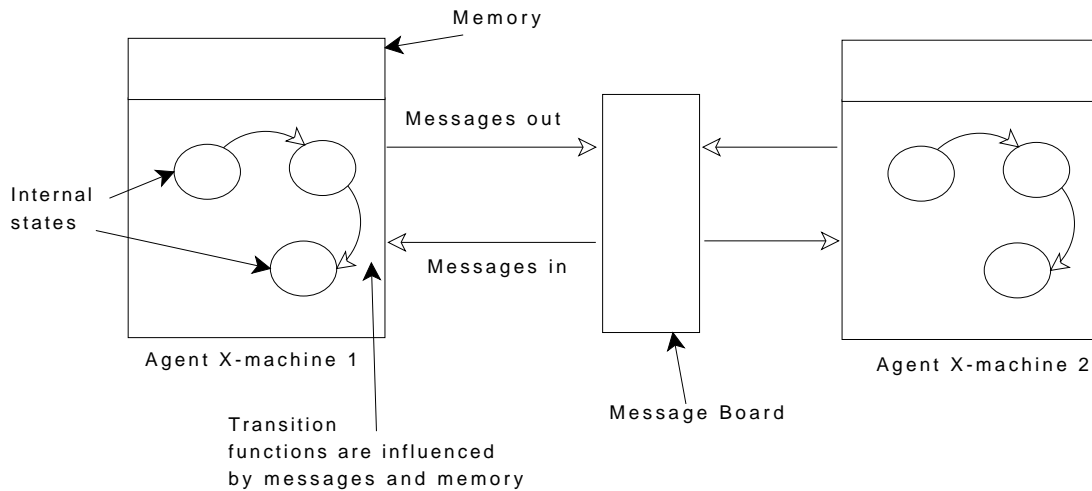


Figure 2: How two agent X-machines communicate. The agents send and read messages from the message board which maintains a database of all the messages sent by the agents.

2.2 Swarm Example

A swarm model in a model which presents the behaviour of birds flocking together. The individual birds follow simple rules, but collectively they produce complex behaviour of the group, as observed in nature. This simple flocking model involves birds to sense where other birds are and then respond accordingly. The activities or functions they perform are:

- Observe if there is a bird nearby.
- Adjust bird position, direction and velocity accordingly.

Converting this model into an agent-based model requires visualising the model as a collection of agents. As the only individuals involved in the model are birds, agents will be representing birds. The functions these bird agents would perform will be:

- Signal. The agent would send information of its current position.
- Observe. The agent would read in the positions from other agents and possibly change velocity.
- Respond. The agent would update position via the current velocity.

The functions would occur in an order as shown in Figure 3. The complete figure represents the functions the agents would be performing during one iteration².

²FLAME prevents the agents to loop back due to parallelization constraints.




Figure 3: Swarm model including states

Figure 4 depicts a situation where their would be conditions added to the functions of the agents. For instance, in the swarm model, there could be a condition added to the z-axis value to determine which response function to perform for the agent. If z is more than zero, the agent would be flying, else if z is zero, then the agent is stationary.




Figure 4: Swarm model including function conditions

The message being used for communication between the agents, in the model, is a signal message, which is the output from ‘signal’ function and the input to the ‘observe’ function (Figure 5). This message includes the position of the agent that sent it with the x , y and z coordinates (Table 1).




Figure 5: Swarm model including messages

An important factor to note here is that FLAME carries the features of a filter which can be added to the messages. This filter can ensure that only the messages in the agents viewing distance are being read, preventing each agent to traverse through all the messages on the message board. The filter will be a formula involving the position contained in the message (the position of the sending agent) and the receiving agent position.

2.3 Transition Functions

Transition functions allow agents to change the state in which they are in, modifying their behaviour. Transition functions take as input the current state s_1 of the agent, current memory value m_1 and the possible arrival of a message that the agent reads t_1 . Depending

Type	Name	Description
double	px	x-axis position
double	py	y-axis position
double	pz	z-axis position

Table 1: Signal Message

on these three values the agent changes to another state s_2 , updates the memory to m_2 and optionally sends a message t_2 .

There could be situations where some of the transition functions do not depend on the incoming messages. Agent transition functions may also be expressed in terms of stochastic rules, which allows the multi-agent systems to be called stochastic systems.

2.4 Memory and States

The difference between the internal set of states and the internal memory set allows the added flexibility when modelling systems. There can be agents with one internal state and all the complexity defined in the memory or equivalently, there could be agents with a trivial memory, with the complexity then bound up in a large state space. It depends on the modeller's perspective on how he/she write the model and where the complexity is added.

In FLAME, one iteration is taken as a standalone run of a simulation. Once all the functions in that iteration have taken place, the message board is emptied, deleting all the messages. This means that messages cannot be sent between iterations, thus models have to be written in a way which considers this.

Table 2 describes the memory variables being used by the bird agents in the swarm model.

Type	Name	Description
double	px	position in x-axis
double	py	position in y-axis
double	pz	position in z-axis
double	vx	velocity in x-axis
double	vy	velocity in y-axis
double	vz	velocity in z-axis

Table 2: Swarm Agent Memory

Modellers can add more variables to the agent memory as they see required. Table 3 represents a transition table presentation of the swarm model. The terms in the table have been defined below:

- Current State - is the state the agent is currently in.
- Input - is any inputs into the transition function.
- M_{pre} - are any preconditions of the memory on the transition.
- Function - is the function name.
- M_{post} - is any change in the agent memory.
- Output - is any outputs from the transition.
- Next State - is the next state that is entered by the agent.

Current State	Input	M_{pre}	Function	M_{post}	Output	Next State
start			signal		signal	1
1	signal		observe	(velocity updated)		2
2		$x > 0$	flying	(position updated)		end
2		$x == 0$	resting	(position updated)		end

Table 3: Swarm Agent Transition Table

The next Section 3 describes how a model can be written up in the XML format that FLAME can understand. Section 7 discusses how to implement the individual agent functions, i.e. M_{post} from the transition table. Section 8 on model execution describes how to use the tools in FLAME to generate a simulation program and execute the simulations.

3 Model Description

Models descriptions are formatted in XML (Extensible Markup Language) tag structures to allow easy human and computer readability. This also allows easier collaborations between the developers writing the application functions that interact with model definitions in the XML.

The DTD (Document Type Definition) of the XML document is currently located at:

<http://eurace.cs.bilgi.edu.tr/XMML.dtd>

For users who are familiar with the HTML structure, a XML document is structured in a similar way as a nested tree structure, where tags contain data or other tags within them. This structure can be condensed into one level or a number of levels within the parent levels. In FLAME, the start and the end of a model file looks like as follows:

```

<?xml version="1.0" encoding="ISO-8859-1"?>
  <!DOCTYPE xmodel SYSTEM
"http://eurace.cs.bilgi.edu.tr/XMML.dtd"> <xmodel version="2">
  <name>Model_name</name>
  <version>the version</version>
<description>a description</description>
  ...
</xmodel>

```

The complete model is contained within the tag level of '*xmodel*'. The *name* of the model is the name of the model being modelled, *version* denotes the version number of the model. The *description* tags allows the model description to be contained in it for modellers to make notes.

3.1 Tags within the XModel

Defining the xmodel is the parent level in the XML file being read by FLAME. This xmodel can be condensed into a number of different tag trees which contain further details about the model. These tags can contain information about:

- **Other models** - Other models can be enabled or disabled when being plugged into a model. This is to allow modeller to test more than one model at a time as well as mix a number of models together.
- **Environment** - The environment contains the global variables of the model in which the agents exist in. Sometimes modellers make the environment act as an agent too with functions and memory states. But this requires another agent to be listed. Here the environment can act as global with constant values for all agents. The environment can contain the following information,
 - Constant variables - Global variables.
 - Location of function files - Location where the functions or C files of the agents are located.
 - Time units - Enables the programming of calendars, which can be assigned to each function to enable it to be active only at specific times during the simulation.
 - Data types - Agent memories can use data structures for some of the variables instead of the traditional C variable types like int, char or double. These data types can be defined by the modeller to contain more than one type or array within it.

- **Agent types** - The agents involved in the system. For instance, in the swarm model, there was only one type of agent the bird agents. In an alternate model of the predator prey model there are two agent types, the fox and the rabbit. These depend on the model being modelled and the modeller's perspectives. The agents are defined by the '*xagent*' tag and can contain the following information,
 - Name - Name of the agent type
 - Description - Textual description of the agent.
 - Memory - A list of the memory variables for each type of agent.
 - Functions - A list of functions the agent can perform. These functions are encapsulated with states like the current and the next state to move to after this function has been executed. The functions would also contain the names of the messages being read in or output from the functions.
- **Message types** - These are a list of all the messages being used in the model. The details within the message are,
 - Name - Name of the message.
 - Description - Textual description of the message.
 - Variables - Variables encapsulated within the message.

Refer to the Appendix to see how these tags are brought together in one model XML file.

3.2 Model in Multiple XML Files

It is possible to define a model in a collection of multiple files. FLAME reads a model from multiple files as if the model was defined in one file. This capability allows different parts of a model to be enabled or disabled easily. For example if a model includes different versions of a sub-model, these can be exchanged, or a subsystem of a model can be disabled to see how it affects the model. Alternatively this capability could be used as a hierarchy, for example a 'body' model could include a model of the 'cardiovascular system' that includes a model of the 'heart'. The following tags show the inclusion of two models, one is enabled and one disabled:

```
<models>
  <model><file>sub_model_1.xml</file><enabled>>true</enabled></model>
  <model><file>sub_model_2.xml</file><enabled>>false</enabled></model>
</models>
```

4 Environment

The environment of a model holds information that maybe required by a model but is not part of an agent or a message. This includes:

- Constant variables - for setting up different simulations easily.
- Location of function files - the path to the implementations of agent functions in C files.
- Time units - for easily activating agent functions dependent on time periods.
- Data types - user defined data types used by agent memory or message variables other than typical C data types.

This notion of environment does not correspond to an environment that would be a part of a model where agents would interact with the environment. Anything that can change in a model must be represented by an agent, therefore if a model includes a changeable environment that agents can interact with, this in itself must be represented by an agent.

4.1 Constant Variables

Constant variables can be set up as part of a simulation for the runs. These are defined as follows:

```
<constants>
  <variable>
    <type>int</type><name>my_constant</name>
    <description>value read in initial simulation settings</description>
  </variable>
</constants>
```

Constant Variables refers to the global values used in the model. These can also be defined in a separate header H file which can then be included in one of the functions C file. The header file should contain the global variable as:

```
#define <varname> <value>
```

This file has to be saved as 'my_header.h' file, include this file into one of the function files so that the compiler knows about these arguments.

4.2 Source C Files

Function files hold the source code for the implementation of the agent functions. These are programmed in C language. They are included in the compilation script (Makefile) of the produced model:

```
<functionFiles>
<file>function_source_code_1.c</file>
<file>function_source_code_2.c</file>
</functionFiles>
```

It is good practise to write source C files in separate files for different agents. But different agent functions should have unique names if they are not using the same function.

4.3 Time Units

Time units are used to define time periods that agent functions act within. For example a model that uses a calendar based time system could take a day to be the smallest time step, i.e. one iteration. Other time units can then use this definition to define other time units, for example weeks, months, and years.

A time unit contains:

- Name - name of the time unit.
- Unit - can contain 'iteration' or other defined time units.
- Period - the length of the time unit using the above units.

An example of a calendar based time unit set up is given below:

```
<timeUnits>
  <timeUnit>
    <name>daily</name>
    <unit>iteration</unit>
    <period>1</period>
  </timeUnit>

  <timeUnit>
    <name>weekly</name>
    <unit>daily</unit>
    <period>5</period>
  </timeUnit>
```



```

<timeUnit>
  <name>monthly</name>
  <unit>weekly</unit>
  <period>4</period>
</timeUnit>

<timeUnit>
  <name>quarterly</name>
  <unit>monthly</unit>
  <period>3</period>
</timeUnit>

<timeUnit>
  <name>yearly</name>
  <unit>monthly</unit>
  <period>12</period>
</timeUnit>

</timeUnits>

```

These time units can be added to the functions, when they are listed as part of the agent. These time units act as conditions on the functions. This has been discussed in Section 5.4.

4.4 Data Types

Data types are user defined data types that can be used in a model. They are a structure for holding variables. Variables can be a:

- Single C fundamental data types - int, float, double, char.
- Static array - of any size for example ten is written as 'variable_name[10]'.
- Dynamic array - available by placing '_array' after the data type name: variable_name_array.
- User defined data type - defined before the current data type.

The example below contains a variable of data structure *position* which contains the x, y and z position in one structure. The position data structure can then be a data type in the *line* data structure.

```

<dataTypes>

  <dataType>
    <name>position/name>
    <description>position in 3D using doubles</description>
    <variables>
      <variable><type>double</type><name>x</name>
        <description>position on x-axis</description>
      </variable>
      <variable><type>double</type><name>y</name>
        <description>position on y-axis</description>
      </variable>
      <variable><type>double</type><name>z</name>
        <description>position on z-axis</description>
      </variable>
    </variables>
  </dataType>

  <dataType>
    <name>line</name>
    <description>a line defined by two points</description>
    <variables>
      <variable><type>position</type><name>start</name>
        <description>start position of the line</description>
      </variable>
      <variable><type>position</type><name>end</name>
        <description>end position of the line</description>
      </variable>
    </variables>
  </dataType>

</dataTypes>

```

5 Agents

A model has to constitute agents. These agents are defined as their type in the model XML file. An agent type contains a name, a description, memory, and functions:

```

<agents>

  <xagent>
    <name>Agent_Name</name>
    <description></description>

```

```

    <memory>
      ...
    </memory>
    <functions>
      ...
    </functions>
  </xagent>

```

5.1 Agent Memory

Agent memory defines variables, where variables are defined by their type, C data types or user defined data types from the environment, a name, and a description:

```

<memory>
  <variable><type>int</type><name>id</name>
    <description>identity number</description>
  </variable>
  <variable><type>double</type><name>x</name>
    <description>position in x-axis</description>
  </variable>
  <variable><type>position</type><name>xyz</name>
    <description>position in x-axis, y-axis, z-axis</description>
  </variable>
</memory>

```

Agent memory variables can be defined as being constant by using the `<constant>` tag and defining it to be true. This will stop the variable being allowed to be changed. This helps message communication in parallel when input filters are dependent upon constant agent memory variables.

```

<variable>
  <type>int</type><name>id</name><constant>true</constant><description></description>
</variable>

```

5.2 Agent Functions

The model XML file requires the agent functions to be listed as well to tell FLAME when the functions will be called in from the C files. An agent function contains:

- Name - the function name which must correspond to an implemented function name
- Description

- Current state - the current state the agent has to be in for this function to execute.
- Next state - the next state the agent will transition to after the function.
- Condition - a possible condition of the function transition.
- Inputs - the possible input messages.
- Outputs - the possible output messages.

And as tags, the XML file will contain:

```

<function>
  <name>function_name</name>
  <description>function description</description>
  <currentState>current_state</currentState>
  <nextState>next_state</nextState>
  <condition>
    ...
  </condition>
  <inputs>
    ...
  </inputs>
  <outputs>
    ...
  </outputs>
</function>

```

The current state and next state tags hold the names of states. This is the only place where states are defined. State names must coordinate with other functions states to produce a transitional graph from the start state to end states.

5.3 Birth and Death of Agents

Agents can be created dynamically by calling the following function in the source code:

```
add_AgentName_agent(var1, var2, ..., var n);
```

The variables are the memory variable which the new agent will have when created. FLAME does not globally track the number of total agents in the environment and thus does not assign unique identifiers to agents. Users can assign their own ids to the agents.

Agents can also die in the environment by returning a 1 in the function. Each agent functions returns a zero. if the function returned a 1, it will be deleted from the environment.

5.4 Function Conditions

A function can have a condition on its transition. This condition can include conditions on the agent memory and also on any time units defined in the environment. Each transition will take the agent from a starting state to an end state at the end of the simulation.

Each possible transition must be mutually exclusive. This means that if a certain condition is true on one part of the branch of functions, there should be an alternate branch which would be the opposite of this condition. This will ensure the model does not halt in the middle during simulation if the condition fails. A function named 'idle' is available to be used for functions that do not require an implementation and a reverse of the conditions.

Conditions (that are not just time unit based) take the form:

- lhs - left hand side of comparison.
- op - the comparison operator.
- rhs - the right hand side of the comparison.

Or in tags form:

```
<lhs></lhs><op></op><rhs></rhs>
```

Sides to compare (lhs or rhs) can be either a value, denoted within value tags or a formula. Values and formulas can include agent variables, which are preceded by 'a'.

```
a.agent_var
```

5.4.1 Comparison Operators

The comparison operator, op, can be one of the following comparison functions:

- EQ - equal to.
- NEQ - not equal to.
- LEQ - less than or equal to.
- GEQ - greater than or equal to.
- LT - less than.
- GT - greater than.

- IN - an integer (in lhs) is a member of an array of integers (in rhs).

Or one of the following logic operators can be used as well:

- AND - for two conditions to be true at the same time.
- OR - for one of the condition to be true.
- NOT - for the condition to be false.

The operator ‘NOT’ can be used by placing ‘not’ tags around a comparison rule. For example the following tagged rule describes the condition being true when the ‘z’ variable of the agent is greater than zero and not greater than ten:

```
<condition>
  <lhs>
    <lhs><value>a.z</value></lhs>
    <op>GT</op>
    <rhs><value>0.0</value></rhs>
  </lhs>
  <op>AND</op>
  <rhs>
    <not>
      <lhs><value>a.z</value></lhs>
      <op>GT</op>
      <rhs><value>10.0</value></rhs>
    </not>
  </rhs>
</condition>
```

Nested conditions can have upto two conditions within them. All conditions have to be mutually exclusive otherwise agents will start disappearing in the resulting xml files.

5.5 Time conditions

A condition can also depend on any time units described in the environment. For example the following condition is true when the agent variable ‘day_of_month_to_act’ is equal to the number of iterations since of the start, the phase, of the ‘monthly’ period, i.e. twenty iterations as defined in the time unit:

```
<condition>
  <time>
```

```

    <period>monthly</period>
    <phase>a.day_of_month_to_act</phase>
  </time>
</condition>

```

The condition allows the function to run *monthly* at the phase of *day_of_month_to_act*.

5.6 Messages in and out of Functions

Functions can have input and output message types. For example, the following example the function takes message types ‘a’ and ‘b’ as inputs and outputs message type ‘c’:

```

<inputs>
  <input><messageName>a</messageName></input>
  <input><messageName>b</messageName></input>
</inputs>
<outputs>
  <output><messageName>c</messageName></output>
</outputs>

```

5.7 Message Filters

Message filters can be applied to message inputs to allow the messages to be filtered. Filters are defined similar to function conditions but use agent memory variables as well as message variables:

```

m.msg_var
a.agent_var

```

The various tags associated with message filters are as follows:

Conditions on the value of a variable within the message. This is denoted by the lhs, op and rhs operators.

The following example filter only allows messages where the agent variable ‘id’ is equal to the message variable ‘worker_id’,

```

<input>
  <messageName>firing</messageName>
  <filter>
    <lhs><value>a.id</value></lhs>

```

```

    <op>EQ</op>
    <rhs><value>m.worker_id</value></rhs>
  </filter>
  <random>>false</random>
</input>

```

The previous example also includes the use of a random tag, set to false, to show that the input does not need to be randomised, as randomising input messages can be computationally expensive. By default all message inputs are not being randomised.

IN tag. Message input filters can now accept the ‘IN’ operator. The IN operator accepts a single integer in the <lhs> tag and an integer array (static or dynamic) in the <rhs> tag. The filter returns true for any single integer that is a member of the integer array. For example:

```

<filter>
  <lhs><value>m.id</value></lhs>
  <op>IN</op>
  <rhs><value>a.id_array</value></rhs>
</filter>

```

The random tag. The random tag defines if the input needs to be randomised or not, either ‘true’ or ‘false’. By default inputs are NOT randomised. If messages need to be processed in a random order to prevent software artifacts, set the random key to be true.

```

  <random>>true</random>

```

If no message randomisation is needed, the random tag does not need to be set since the default is already set to false.

The sort tag. A sort can be defined for a message input by defining the message variable to be sorted, the ‘key’, and the order of the sort, either ‘ascend’ or ‘descend’. The following example orders the messages in descending order for the values of the variable ‘wage’. If both random and sort are used, the sorting takes precedence over randomisation. After the sort, equally ranked messages are being randomised.

```

  <sort><key>wage</key><order>descend</order></sort>

```

The message filters operate on a deeper level by pre-sorting messages on the message board by the Message Board Library (libmboard).

6 Messages

Messages defined in a model must have a type which is defined by a name and the variables that are included in the message. The following example is a message called ‘signal’ that holds a position in 3D.

```
<messages>

<message>
  <name>signal</name>
  <description>Holds the position of the sending agent</description>
  <variables>
    <variable><type>double</type><name>x</name>
      <description>The x-axis position</description>
    </variable>
    <variable><type>double</type><name>y</name>
      <description>The y-axis position</description>
    </variable>
    <variable><type>double</type><name>z</name>
      <description>The z-axis position</description>
    </variable>
  </variables>
</message>

</messages>
```

7 Model Implementation

The implementations of each agent’s functions are currently written in separate files in C language, suffixed with ‘.c’. Each file must include two header files, one for the overall framework and one for the particular agent that the functions are for. Functions for different agents cannot be contained in the same file. Thus, at the top of each file two headers are required:

```
#include "header.h"
#include "<agentname>_agent_header.h"
```

Where ‘<agent_name>’ is replaced with the actual agent type name. Agent functions can then be written in the following style:

```
/*
 * \fn: int function_name()
```

```

    * \brief: A brief description of the function.
    */
int function_name() {
    /* Function code here */

    return 0; /* Returning zero means the agent is not removed */
}

```

The first commented part (four lines) is good practice and can be used to auto-generate source code documentation. The function name should coordinate with the agent function name and the function should return an integer. The functions have no parameters. Returning zero means the agent is not removed from the simulation, and one removes the agent immediately from the simulation.

7.1 Accessing Agent Memory Variables

After including the specific agent header, the variables in the agent memory can be accessed by capitalising the variable name:

```
AGENT_VARIABLE
```

To access elements of a static array use square brackets and the index number:

```
MY_STATIC_ARRAY[index]
```

To access the elements and the size of dynamic array variables use `‘.size’` and `‘.array[index]’`:

```
MY_DYNAMIC_ARRAY.size MY_DYNAMIC_ARRAY.array[index]
```

To access variables of a model data type use `‘.variablename’`:

```
MY_DATA_TYPE.variablename
```

7.1.1 Using Model Data Types

The following is an example of how to use a data type called *vacancy* which was defined in the model XML file:

```

/* To allocate a local data type */
vacancy vac;

/* And initialise */
init_vacancy(&vac);

/* Initialise a static array of the data type */
init_vacancy_static_array(&vac_static_array, array_size);

/* Free a data type */
free_vacancy(&vac);

/* Free a static array of a data type */
free_vacancy_static_array(&vac_static_array, array_size);

/* Copy a data type */
copy_vacancy(&vac_from, &vac_to);

/* Copy a static array of a data type */
copy_vacancy_static_array(&vac_static_array_from,
                          &vac_static_array_to, array_size);

```

If the data type is a variable from the agent memory, then the data type variable name must be capitalised.

7.1.2 Using Dynamic Arrays

Dynamic array variables are created by adding ‘_array’ to the variable type. The following is an example of how to use a dynamic array:

```

/* Allocate local dynamic array */
vacancy_array vacancy_list;

/* And initialise */
init_vacancy_array(&vacancy_list);

/* Reset a dynamic array */
reset_vacancy_array(&vacancy_list);

/* Free a dynamic array */
free_vacancy_array(&vacancy_list);

/* Add an element to the dynamic array */
add_vacancy(&vacancy_list,

```

```

var1, .. varN);

/* Remove an element at index */
remove_vacancy(&vacancy_list,
index);

/* Copy the array */
copy_vacancy_array(&from_list, &to_list);

```

If the dynamic array is a variable from the agent memory, then the dynamic array variable name must be capitalised.

7.2 Sending and receiving messages

Messages can be traversed with in a function. The messages can be read using macros to loop through the incoming message list as per the template below, where ‘messagename’ is replaced by the actual message name. Message variables can be accessed using an arrow ‘->’:

```

START_MESSAGENAME_MESSAGE_LOOP
messagename_message->variablename
FINISH_MESSAGENAME_MESSAGE_LOOP

```

Messages are sent or added to the message list by,

```

add_messagename_message(var1, .. varN);

```

8 Model Execution

FLAME contains a parser program called ‘xparser’ that parses a model XML definition into simulation program source code. This can be compiled together with model implementation source code for the simulations. The xparser includes template files which are used to generate the simulation program source code.

The xparser takes as parameters the location of the model file and an option for serial or parallel (MPI) version, serial being the default if the option is not specified.

8.1 Xparser Generated Files

The Xparser generates simulation source code files in the same directory as the model file. These files are:

- Doxyfile - a configuration file for generating documentation using the program ‘doxygen’.
- header.h - a C header file for global variables and function declarations between source code files.
- low_primes.h - holds data used for partitioning agents.
- main.c - the source code file containing the main program loop.
- Makefile - the compilation script used by the program ‘make’.
- memory.c - the source code file that handles the memory requirements of the simulation.
- xml.c - the source code file that handles inputs and outputs of the simulation.
- <agent_name>_agent_header.h - the header file containing macros for accessing agent memory variables.
- rules.c - the source code file containing the generated rules for function conditions and message input filters.
- messageboards.c - This is done automatically now.
- partitioning.c - Not used anymore for partitioning. This is automatically done.
- latex.c - generates a latex file containing the agent names and variables defined as tables for report writing.

For running in parallel, additional files are generated:

- propagate_messages.c - deprecated?
- propagate_agents.c - still used?

The simulation source code files then require compilation, which can be easily achieved using the included compilation script ‘Makefile’ using the ‘make’ build automation tool. The program ‘make’ invokes the ‘gcc’ C compiler, which are both free and available on various operating systems. If the parallel version of the simulation was specified the compiler invoked by ‘make’ is ‘mpicc’ which is a script usually available on parallel systems.

The compiled program is called ‘main’. The parameters required to run a simulation include the number of iterations to run for and the initial start states (memory) of the agents, currently a formatted XML file.

8.2 Start States Files 0.xml

The format of the initial start states XML is given by the following example:

```
<states>
  <itno>0</itno>
  <environment> <my_constant>6</my_constant>
</environment>
  <xagent>
  <name>agent_name</name>
  <var_name>0</var_name>
  ...
</xagent>
  ...
</states>
```

An example of this File has been presented below:

```
<states>
  <itno>0</itno>
  <xagent>
  <name>Firm</name>
  <id>1</id>
  <age>0</age>
  <region_id>1</region_id>
  <bank_id>1233</bank_id>
  <gov_id>1232</gov_id>
  <employees>{}</employees>
  <wage_offer>1.000000</wage_offer>
  <wage_offer_for_skill_1>1.000000</wage_offer_for_skill_1>
  <wage_offer_for_skill_2>1.000000</wage_offer_for_skill_2>
  <wage_offer_for_skill_3>1.000000</wage_offer_for_skill_3>
  <day_of_month_to_act>0</day_of_month_to_act>
  <posx>0.000000</posx>
  <posy>0.000000</posy>
  </xagent>
</states>
```

The root tag is called 'states' and the 'itno' tag holds the iteration number that these states refer to. If there are any environment constants these are placed within the 'environment' tags. Any agents that exist are defined within 'xagent' tags and require the name of the agent within 'name' tags. Any agent memory variable (or environment constant) value is defined within tags with the name of the variable. Arrays and data types are defined within curly brackets with commas between each element.

It is not needed to copy the entire 0.xml file into the folder hierarchy. Instead, you could just have a ‘shell’ xml file that specifies the output tags and constants and then refers to the big 0.xml file and imports the agent memory settings from outside, like

```
<states>
  <outputs>
    <output>
      <type>snapshot</type><name/>
      <location>./</location>
      <format>xml</format>
      <time><period>501</period><phase>0</phase><duration/></time>
    </output>
    <output><type>agent</type><name>Eurostat</name><location>./</location><format>xml</format>
      <time><period>20</period><phase>0</phase><duration/></time>
    </output>
  </outputs>

  <imports>
    <import>
      <location>../../0_Integrated_2R_default.xml</location>      # the location of the data
      <format>xml</format>      # the format of the data
      <type>environment</type>      # environment or agent data
    </import>
    <import>
      <location>../../0_Integrated_2R_default.xml</location>      # the location of the data
      <format>xml</format>      # the format of the data
      <type>agent</type>      # environment or agent data
    </import>
  </imports>
</states>
```

When a simulation is running after every iteration, a states file is produced in the same directory and in the same format as the start states file with the values of each agent’s memory.

9 Installation Notes on FLAME

FLAME required a number of softwares installed on the computer to be able to execute. The recommended C compiler used in the C compiler is gcc from MSys1.0 and make.exe from MinGW. These are:

- Latest version of the framework 1.0 (Xparser)

- C compiler (MinGW)
- Libmboard

9.1 MinGW

Recommended C compiler is gcc which is built-in Unix. For windows, instead use the make.exe from msys.

Download site for Msys 1.0:

<http://www.mingw.org/wiki/MSYS>

Configuring MinGW for Windows users, follow the following path:

```
Computer -> Properties -> Advanced settings -> Environment variables
-> System variables
```

Select 'Path' and edit it as follows:

- Add the path to bin folder: C:\msys\1.0\bin.

9.2 Libmboard

For Windows: Download the libmboard for windows and unzip and place the folder where your model is. This is an already precompiled version for windows platforms.

For Linux/Mac systems:

1. Download the latest version of libmboard from ccpforge. Place this anywhere because we will compile this and put it to a specific place on the operating system to be used with FLAME.
2. Go to the Folder where you want to place libmboard. For example if placing on '\Volumes':
 - mkdir libmboard. This make a directory for libmboard.
3. Go to the downloaded libmboard. Unzip this and go into the folder.

```
>./configure --prefix =/Volumes/libmboard --disable -tests
>make
```


If no complaints, go into the folder ‘/Volumes/libmboard’ and run make install as superuser:

```
>sudo make install
```

This will compile the libmboard on your system.

9.3 Extra Useful Installations

9.3.1 GDB

For debugging, GDB GNU Debugger is recommended. It is freely available with many tutorials on how to use on web. You can get your free copy from the following link:

http://sourceforge.net/project/showfiles.php?group_id=2435&package_id=20507

Note: Windows users are recommended to use the version 5.2.1 (available in the above link).

9.3.2 Dotty

The parser creates diagrams about the flow of your model. These are created in ‘.dot’ format. To access these files download Graphviz from www.graphviz.org. You can view and include these pictures in your model description and convert the image to other formats like pdf.

10 Running a Simulation

After writing the model XML file and C functions files of the agent, the xparser has to be used to compile the simulation program. This is done by going into where the xparser is placed and writing the following commands:

```
FLAME_xparser> xparser.exe ../model/model.xml
```

This creates all files which contain details of running the program. Extra files are created in ‘.dot’ format which can be opened using Graphviz. The dot files represent graph structures of the agents which show a description of how the model will work.

By default, the xparser will generate files for running the model in a serial format. If parallel version of the model was required then just an extra tag has to be added,

```
FLAME\_xparser> xparser.exe ../model/model.xml -p
```

The parallel version, by default, produces code for geometric partitioning of the agents depending on the locations.

After creating these files, users have to go into the folder where the model was located and compile the files.

```
Model>make
```

This creates a main program which is the main simulation program. The main.exe file can then be linked with the initial start states and the number of iterations wanted to be written out.

```
Model>main.exe 10 its/0.xml
```

Main.exe is the simulation program, 10 is the number of iterations to produce and its/0.xml is the initial start states of the model which the modeller defined.

The default variables used for geometric partitioning are as follows:

- x , y, z
- posx, posy, posz

If the code has to partitioning in a round robin manner the flag used is -r. Example of this as follows,

```
Model>mpirun -np 2 main.exe 10 its/0.xml -r
```

If the model is being executed in parallel, the mpirun is called to use MPI (Message passing Interface) for running the model. 2 denotes the number of nodes the model is being divided over and the '-r' flag denotes a round robin distribution of the agents over the nodes. This flag is optional.

For Mac or Linux users, 'main.exe' files are written as './main'. This is true for all exe files run on a similar platform.

10.1 Running Job Scripts

This example is given with reference to running job scripts on a parallel machine, Iceberg. The high performance computing grid Iceberg had to be configured to use FLAME. The steps involved were as follows,

1. Connection to Iceberg via a username and password.
2. Copy files to Iceberg which includes xparser and libmboard files.
3. Copy files from Iceberg which includes the resulting xml files.
4. Configuring the C++ compiler and MPI libraries on Iceberg. This was done by creating a symbolic link to all the required files like cc1plus.
5. Configuring the libmboard by running ./configure. The MPI libraries are linked in here.
6. Run a model on Iceberg in parallel by using the following commands:

```
./xparser path_of_model.xml -p
```

For example,

```
./xparser ../model/turningKernel.xml -p
```

```
make CC=/usr/local/packages5/openmpi-gnu/bin/mpicc
LIBMBOARD_DIR=/home/ac1mk/libmboard export
LD_LIBRARY_PATH=\$LD_LIBRARY_PATH:/usr/local/packages5/openmpi-gnu/lib
```

Then run the parallel nodes by submitting a job.sh file which contains the number of nodes and the time you specify to it.

```
#!/bin/sh

#\$ -l h_cpu=0:30:00 (Specifies how long the job is likely to last)

#\$ -cwd #\$ -pe openmpi-ib 8 (Specifies the openmpi-ib library should
be used, with 8 cores)

#\$ -q parallel.q (Specifies the queue to be used)

#\$ -v SGE_HOME=/usr/local/sge6_2 (Specifies the path)

/usr/mpi/gcc/openmpi-1.2.8/bin/mpirun /home/ac1mk/trial2/main 100
trial2/output/p8/t100k/0.xml -r
```

On a mac system this could be done directly by,

```
mpirun -np numberOfNodes ./main numberOfIterations 0.xml -r
```

Such as,

```
mpirun -np 16 ./main 100 0.xml -r
```

11 Notes for the Programmers of FLAME

11.1 Memory Allocation and its Problems

Memory allocation for the agents and the messages is done as a continuous block size of memory. The command *sizeof* is used to return a byte size of the agent memory in use. This is an important facet for parallelization when using MPI. Sending data from one node to the other requires the program to know how many bytes have to be sent across to package it up in small packets. Thus it becomes important to determine its size.

11.1.1 Dynamic Arrays

FLAME also allows the use of dynamic arrays which causes a hindrance to this area of parallelization. It is strongly discouraged for dynamic arrays to be used as part of the agent memory, if the agents have to be moved around in parallel. Dynamic arrays also prevents the associated memory to be allocated as blocks of continuous memory. Messages are another factor which discourages the use of dynamic arrays within the messages. The size of the message becomes difficult to be determined and sent to and fro for this reason.

11.1.2 Data Types

User-defined data types are allocated as pointers in agent memory but this has been modified in a new version to be released. This means that instead of user using an arrow ‘->’ to dereference variables, a dot ‘.’ is used to access the data structure.

Dynamic array data structures are also not allocated as a pointer (but the actual dynamic array is) which means functions to interact with a dynamic array data structure need to pass a pointer. This means the use of the ampersand ‘&’ to reference the data structure.

11.1.3 Agent Memory Management

Each agent has an associated memory data structure. Since the early versions of the framework all agents have been managed in one list. This was so that the list could be randomised and therefore remove any chances of agents having priority over other agents by always being executed first. In essence, the same effect can be achieved by randomising the messages output and therefore the message inputs into agents. The current framework has a generic agent memory structure that can point to any specific agent type.

With the introduction of the new message board library the action of randomising (or now also sorting and filtering) messages the need to randomise the agent list is redundant. Also redundant is the need to have a single list of all the agents. The generic agent memory structure is therefore not needed and each agent type can have it's only separate list.

11.2 Agent Execution

Agents have a number of functions to perform. The order that these functions are run is defined by the states associated with each function.

Figure 6 depicts an example of how the states can link functions together. In the first case, the agent performs two functions during the iteration step. The current state and the next state determine the order of the functions. *Function A* is followed by *Function B* by simply assigning the current and next states to link the function chain together. Case 2, presents another scenario, where the *Function A* is run twice during a simulation step. The same function can be run twice by linking it to different current and next states.

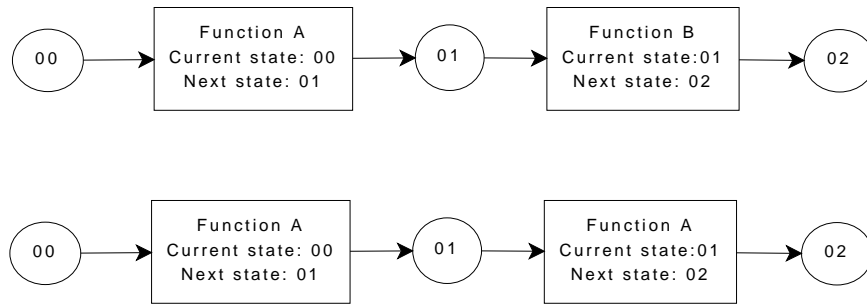


Figure 6: Using states to form a function chain during one iteration step.

These order of states also determines the internal dependency between the functions. This is only true if only one agent is being discussed. But if there is a dependency between more than one agent, communication dependencies are generated. These are denoted by the messages being sent and read by other functions.

The communication dependency sets up a synchronisation point as shown in Figure 7. This means that all agent *As* have to finish running their *Function A* before they can start running the *Function D* for agent *B*.

Using the X-machine methodology, the agents traverse through the states to run the defined functions. These functions are also the transition functions which are defined in the model XML file with the,

- current state: the current state of the agent
- input: the inputs the function is expecting
- m_{pre} : the conditions on memory of executing the function
- name: the name of the function
- m_{post} : the changes in the memory (i.e. the function code)

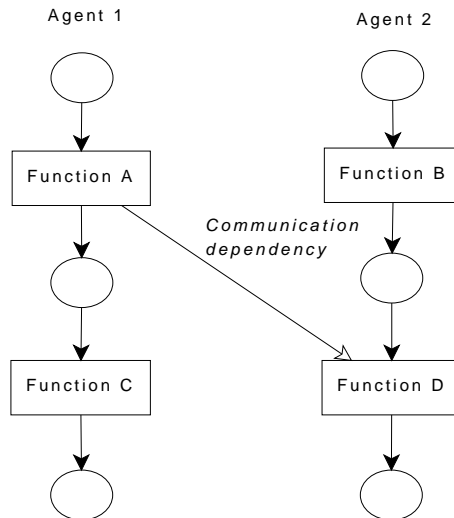


Figure 7: Function D of Agent 2 depends on all Agent 1s to finish their Function A.

- output: possible outputs of the function
- next state: the next state to move the agent to

By producing an order of function execution, this also provides a way to manage the processing of agents. By providing a link to an agent list for each possible agent state, agents can be moved between these agent state lists until they reach an end state.

11.3 Communication

MPI or Message Passing Interface is used to handle the communication between the agents. Using MPI has a number of advantages:

- MPI allows language independent communication, which means that different platforms can be linked together to communicate messages.
- Synchronisation between message channels.
- Shared or distributed memory.
- Packaging messages into blocks of memory to be sent across.

11.4 Libmboard

The communication handling was decoupled with the FLAME implementation and programmed as a separate message library. This was done to provide more flexibility with different parallelisation strategies to the current FLAME modellers.

The Message board library defines a set of routines and functions which can be integrated with the FLAME code to parse messages. This was called the libmboard and uses MPI to communicate between processors. Details of the Message Board can be found in its Reference Manual at

<http://www.softeng.cse.clrc.ac.uk/libmboard>

12 XParser Distribution

The xparser is distributed as a series of template files accompanied with a few header files. These template files can be downloaded in to the desired directory. The freely available GCC compiler is then used to compile the files on the machine.

The libmboard (or the Message Board) is an additional feature of the xparser which is being developed to increase the efficiency of parallel communication of large computers. This file can also be downloaded and compiled on the machine for running the simulations.

13 XParser generated files for FLAME programmers

Reading the accompanying document '*FLAME User Manual*', it is explained that when executing the xparser with the model, a number of files are generated as part of the simulation package. These files are as follows,

- Doxyfile - Generated documentation for the model.
- Header files for each agent memory - Contains pointers for accessing agent memory during simulation.
- Header.h - Memory for the xparser.
- Low_primes.h - For partitioning of the agents.
- Main.c - The main C code for running the simulation.
- Main.exe - The simulation file.

- MakeFile - Makefile contains the details of the locations of the files, flags associated etc.
- Memory.c - contains the memory functions like reading through messages or agents.
- Messageboard.c - deprecated. Not needed any more as automatically done.
- Partitioning.c - deprecated. Not needed any more as automatically done.
- Rules.c - deprecated. Not needed any more as automatically done.
- Xml.c - Contains functions to parse through the XML file.

Figure 8 describes a series of steps which the xparser goes through to generate a simulation package of the model. These steps have been explained as follows:

1. Reading in the model. The template Readmodel.tmpl provides these functions. This file allows reading of the various tags in the model XML file.
2. Creating of dependency graph. The dot files are generated which are known as a series of stategraph files. These diagrams display the description of the model, which order the functions are called in and the different layers which denote the synchronisation points among the agent functions due to communication dependencies.
3. Writing out the make file. This file contains the location of various files, like the libmboard and more.
4. Writing out the xml.c file. This file contains functions for reading specific data variables like static arrays, ints or doubles with in the agent memory. It also contain functions for reading the data structures within the agent memory, for example:

```
read_mall_strategy(char * buffer, int * counter, mall_strategy *
tempdatatype)
```

The file also contains functions on reading the initial starting states file for the agent memories. For this purpose it opens the '0.xml' file and reads these values.

For parallel computation, an array is initialised for allowing round robin distribution of the agents.

5. Writing out the main.c file. This is the main file which contains the complete xparser functions being called. It reads in the number of iterations needed for the simulation, the initial start states file, generate partitions for parallel computing and saves the iteration data in progressive XML files. This file also performs additional functions, like

- traverses through the different agent states.

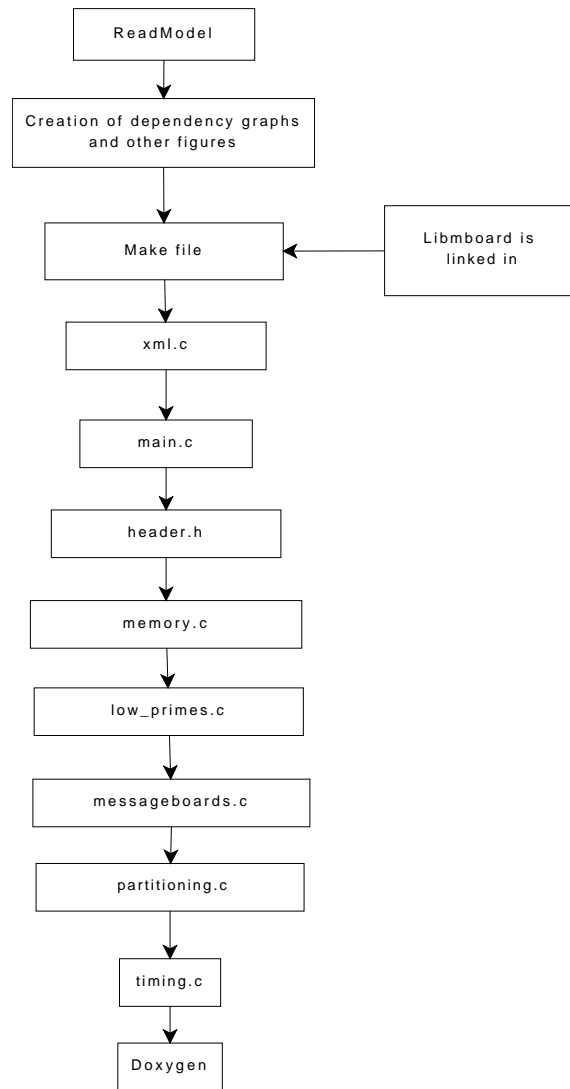


Figure 8: Block diagram of the series of files read for creating the model simulation package.

- checks conditions of the agent functions before calling them.
- calls the synchronisation code for the message boards. These are specific MB functions which have been documented in the libmboard documentation.
- creates iterator for the messages.
- freeing agents when moving to next states.
- clears the message board at the end of the iteration.
- clean up.

6. Header.h file. This file provides the names being used in the simulation. For instance, the xmachine agent memory, the states and the prototypes for the agent functions.

Various definitions for the messageboards like names of the iterators are also defined here, along with prototypes for reading and writing various agent memory variables.

7. Writing the memory.c file. Memory.c file contains the actual function code of the functions being used by the xparser. These are functions like free agents, or freeing messages by calling MB_Clear. The file contains various functions to initialise memory variables like arrays or data structures. Functions for adding and freeing agents are also contained here.

8. Writing the rules.c file. Rules .c contains the rules being used in the model. For instance,

- For function conditions,

```
iteration loop%20==6 return 1 else return 0;
```

- Or for agent memory conditions,

```
a->learningwindow==0 return 1 else return 0;
```

9. Writing the low_primes.c file. Low primes file defines the arrays which are used for partitioning of the data.

10. Writing the messageboards.c files. Messageboard.c is used for writing functions which allows access to the messageboard. For instance,

```
/*for adding messages*/  
MB_AddMessage(b_messagename,&msg)  
  
/*Rewinding an iterator*/ MB_Iterator_Rewind(i_mall_strategy_to_use)  
  
/* getting a message*/  
MB_Iterator_GetMessage(i_mall_strategy_to_use, (void**)&msg);
```

11. Writing the partitioning.c file. Partitioning.c file contains details for generating partitions as geometric nodes and saves this data.

12. The timing.c returns the time it takes to run the code.

13. The Doxygen file writes out data about the model file.

Details of the message board functions are available at www.softeng.cse.clrc.ac.uk/wiki/EURACE/MessageBoards/ImplementationNotes/API

14 XParser Versions

During the development process, the Xparser has gone through a series of development versions, each being modified to include more features for making use easy for the modellers and increasing efficiency.

Change logs for the different versions has been stored at the CcpForge repository for the developers. The XParser has a number of versions, with the latest version 0.16.2.

The current final version xparser 0.16.2 has been tested and used for various simulations of the economic EURACE model and has been proved to be very stable and good for economic modelling.

This version has thus been tagged to be released as FLAME version 1.0. Any future updates to this version will be announced as new versions of FLAME.

15 Example Models

FLAME has a large community of modellers who are using the software for their own research. Some of these models have been packaged up and distributed as example models to help future FLAME users in doing their research. There are freely distributed with the XParser.

A XML DTD

```
<!ELEMENT xmodel
  (name,version,description,models?,environment?,agents,messages?)>
<!ATTLIST xmodel version CDATA #REQUIRED>

<!ELEMENT models (model*)>
<!ELEMENT model (file,enabled)>

<!ELEMENT environment (constants?,functionFiles?,timeUnits?,dataTypes?)>
<!ELEMENT dataTypes (dataType*)>
<!ELEMENT dataType (name,description,variables)>
<!ELEMENT variables (variable*)>
<!ELEMENT variable (type,name,description)>
<!ELEMENT constants (variable*)>
<!ELEMENT functionFiles (file*)>
<!ELEMENT timeUnits (timeUnit*)>
<!ELEMENT timeUnit (name,unit,period)>
```

```

<!ELEMENT agents (xagent*)>
<!ELEMENT xagent (name,description,memory?,roles?,functions?)>
<!ELEMENT memory (variable*)>
<!ELEMENT functions (function*)>
<!ELEMENT function
  (name,description,code?,currentState,nextState,condition?,inputs?,outputs?)>
<!ELEMENT condition ((not)|(lhs,op,rhs)|(time))>
<!ELEMENT not ((lhs,op,rhs)|(time))>
<!ELEMENT lhs ((not)|(lhs,op,rhs)|(value)|(time))>
<!ELEMENT rhs ((not)|(lhs,op,rhs)|(value)|(time))>
<!ELEMENT time (period,phase,duration?)>
<!ELEMENT inputs (input*)>
<!ELEMENT input (messageName,filter?,sort?, random?)>
<!ELEMENT filter (lhs,op,rhs)>
<!ELEMENT outputs (output*)>
<!ELEMENT output (messageName)>
<!ELEMENT messages (message*)>
<!ELEMENT message (name,description,variables)>

<!ELEMENT code (#PCDATA)>
<!ELEMENT currentState (#PCDATA)>
<!ELEMENT description (#PCDATA)>
<!ELEMENT enabled (#PCDATA)>
<!ELEMENT file (#PCDATA)>
<!ELEMENT messageName (#PCDATA)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT nextState (#PCDATA)>
<!ELEMENT op (#PCDATA)>
<!ELEMENT sort (#PCDATA)>
<!ELEMENT random (#PCDATA)>
<!ELEMENT statement (#PCDATA)>
<!ELEMENT type (#PCDATA)>
<!ELEMENT value (#PCDATA)>
<!ELEMENT version (#PCDATA)>
<!ELEMENT unit (#PCDATA)>
<!ELEMENT period (#PCDATA)>
<!ELEMENT phase (#PCDATA)>
<!ELEMENT duration (#PCDATA)>

```