

CoHLA – User Manual

<https://cohla.nl/>

Thomas Nägele
Radboud University
thomas@cohla.nl

2nd August 2019

Contents

1	Introduction	2
2	Installation	2
3	Definitions	2
4	CoHLA project definition	2
4.1	Imports	3
4.2	Environment (required)	3
4.2.1	HLA implementation (required)	3
4.2.2	Source directory	3
4.2.3	Print level	3
4.2.4	Publish only changes	4
4.2.5	Example	4
4.3	Federate classes (required)	4
4.3.1	Type (required)	4
4.3.2	Type configuration	4
4.3.3	Attributes	5
4.3.4	Parameters	6
4.3.5	Time policy	6
4.3.6	Default model	6
4.3.7	Advance type	6
4.3.8	Default step size	7
4.3.9	Default lookahead	7
4.3.10	Simulation weight	7
4.4	ConnectionSets	7
4.4.1	ConnectionSet connection	7
4.5	Configurations	7
4.5.1	Initialisation attributes	8
4.6	Federations	8
4.6.1	Instances	8
4.6.2	Connections	8
4.6.3	Situations	9
4.6.4	Scenarios	9
4.6.5	Fault scenarios	10
4.6.6	Design Space Exploration	10
4.6.7	Metric sets	11
4.6.8	Distributions	12

5	Running a co-simulation	12
5.1	Configuration files	12
5.2	Run-script	12
5.2.1	Basic behaviour	13
5.2.2	Configurables	13
5.2.3	Distribution setup	13
5.2.4	Other options	14
5.2.5	Federate-specific options	14
5.2.6	Useful commands	14
A	CoHLA grammar	15

1 Introduction

This manual describes the basics on how to use CoHLA (Configuring HLA). The project aims for an easier implementation of co-simulations of systems using the High Level Architecture. Using a DSL, one can easily describe a model-based co-simulation. From this co-simulation definition, the framework will generate code that can be compiled and simulated using OpenRTI.

2 Installation

This manual assumes that CoHLA and all of its dependencies has already been installed on the system. For the installation, there are two options.

- Install CoHLA and its dependencies on your system using the Installation Manual¹.
- Download and import the VirtualBox² image to use the fully installed and configured Ubuntu 18.04 system. The image is available on <https://cohla.nl/downloads/>.

3 Definitions

Below we introduce the terminology used in this manual.

CoHLA Workspace This is the workspace where the CoHLA projects will be created and maintained.

Project Directory This is the directory in the CoHLA Workspace where the project is located. This directory may contain the models (possibly in a sub-directory) and all CoHLA source files.

Source Generation Directory This is the directory in the Project Directory where the sources generated by the CoHLA source generator will be created. These sources are typically located in the *src-gen* directory in the Project Directory. This directory is overwritten every time the generation process is invoked. Hence, permanent files should not be placed here.

Model Directory This may be a directory in the Project Directory in which all models for the co-simulation are stored. The Model Directory may also be equal to the Project Directory.

Library Directory The directory in which the OpenRTI libraries are installed. By default, these libraries are installed in `C:\opt\OpenRTI-libs` (for Windows) or `/opt/OpenRTI-libs` (for Linux/Mac).

4 CoHLA project definition

This section outlines all ingredients that can be used to specify a co-simulation of models using CoHLA. The subsection layout of this section corresponds to the order in which the components should be defined. Sometimes, an example is shown to demonstrate the use of a set of properties that was introduced earlier. These subsections are typically named “Example”.

¹<https://cohla.nl/docs/>

²<https://www.virtualbox.org/>

4.1 Imports

CoHLA allows the use of imported CoHLA files. This can be particularly useful when specifying Federate Classes in different source files. An import consists of the keyword `import` followed by the file name to import as a string. Zero or more imports may be specified in a file.

```
1 | import "../models/model.cohla"
```

4.2 Environment (required)

An environment block is used to specify some code generation settings.

```
1 | Environment { ... }
```

4.2.1 HLA implementation (required)

Specifies the HLA RTI implementation that is used. Currently, the generator is only capable of generating code for OpenRTI. The HLA implementation block starts with the keyword `RTI`. The block contents starts with the RTI implementation to generate code for, which is one of the following RTIs.

`OpenRTI` Generates C++ code for OpenRTI³.

`PitchRTI` Generates C++ or Java code for Pitch⁴.

`Portico` Generates Java code for Portico⁵.

The file path to the libraries for this RTI should also be specified by using the `Libraries` keyword, followed by a string containing the path. Dependencies may be located on a different location on the file system. This location may be provided by using the optional `Dependencies` keyword.

```
1 | RTI {  
2 |   OpenRTI  
3 |   Libraries "/opt/orti-libs"  
4 |   Dependencies "/opt/deps" // Optional  
5 | }
```

4.2.2 Source directory

Overrides the default output directory for the sources that are being generated. The directory is set to "src" by default.

```
1 | SourceDirectory "sources"
```

4.2.3 Print level

Sets the print level of each of the generated federates. By default, on every advancement in time, the state of the federate (its attribute values) are being printed to the console. Other options are the following.

`State` Prints all its attribute values upon time advancement. (default)

`Time` Prints only the time upon time advancement.

`None` Disable printing.

```
1 | PrintLevel None
```

³<https://sourceforge.net/projects/openrti/>

⁴<http://www.pitchtechnologies.com/products/prti/>

⁵<http://www.porticoproject.org/>

4.2.4 Publish only changes

When the keyword `PublishOnlyChanges` is specified, all federates only publish attributes after time advancement when one or more of its published attributes has changed.

4.2.5 Example

A sample definition of an Environment is displayed below.

```
1 | Environment {
2 |   RTI {
3 |     OpenRTI
4 |     Libraries "/opt/OpenRTI-libs"
5 |   }
6 |   PublishOnlyChanges
7 | }
```

4.3 Federate classes (required)

A Federate Class represents a simulation model that could be connected to the co-simulation. Multiple instances of the same Federate Class can be connected to the co-simulation. The Federate Class is used to specify some default parameters and its attributes. Every Federate Class is identified by a name.

```
1 | FederateClass SimModel { ... }
```

4.3.1 Type (required)

To specify the type to use for the model, a `Type` must be specified. Values can be one of the following.

- **POOSL**: POOSL model using the Rotalumis simulator.
- **FMU**: FMI model to be controlled using the FMI=library.
- **CSV-logger**: Logger federate exporting a CSV file.
- **BulletCollision**: Collision detection using the Bullet Physics Library.
- **None**: No simulator specified.

When `None` is specified, no wrapper code for a simulator is generated. A sample definition is the following.

```
1 | SimulatorType FMU
```

4.3.2 Type configuration

Some simulator types require some additional configuration to work. For now, this is only the case for POOSL and CSV-logger federates. The configuration is specified in a block, directly following the simulator type.

POOSL configuration

Processes are used to specify a comprehensive mapping to attributes in a POOSL model. Since a POOSL model executes as a set of processes, a path to the process in which an attribute is defined is required to read the attribute from the simulator. The identifiers of these processes are then used when defining attributes for the Federate Class. All processes should be specified within a `Processes` block.

```
1 | Processes { ... }
```

Every process consists of a name (ID) and a path (string) , separated by the keyword `in`:
`{ID} in {STRING}`.

```
1 | Processes {
2 |   process1 in "path/to/process"
3 |   process2 in "path/to/second/process"
4 | }
```

CSV-logger configuration

The optional configuration for a CSV-logger type of federate only consist of a default measure time to record the simulation. This value is provided after the `DefaultMeasureTime` keyword.

```
1 | DefaultMeasureTime 5800.0
```

4.3.3 Attributes

Attributes are variables that may be shared with other simulators or act as inputs from other simulators (or both). Basic attributes can be defined within a single line, although some more HLA-specific configuration can be done in an option block. Attributes must be defined in an `Attributes` block.

```
1 | Attributes { ... }
```

A basic attribute is defined by the following:

```
{SharingType} {DataType} [Collision]? [{MI_Operator}]? {ID} {Alias}?
```

The `Collision` keyword sets the attribute to be the receiver of collider state for the collision simulator. When the `DataType` is Boolean, this is either `true` or `false`; when it is an Integer, the number of collisions found is stored.

SharingType

The sharing type specifies whether the attribute is an input or an output attribute. The attribute can also be both input and output. Possible values are `Input`, `Output`, `InOutput`.

DataType

The data type specifies the data type the attribute stores. It can either be `Integer`, `Long`, `Real`, `Boolean` or `String`, all corresponding to their C++ or Java native types.

MI_Operator

The MultiInput operator specifies how to determine the attribute value when multiple outputs are connected to it. For example, an input attribute that receives values from three different sources i_1 , i_2 and i_3 combines the latest attribute values of these sources into one value for the attribute. The latest values of each of the sources are cached to be able to recompute the correct value of the attribute. By default the `MI_Operator` is `None`, which means that the latest value that is received is stored, thus overwriting the older value, which is possibly from another source. Possible values and their meaning are listed below.

None (default) No operator specified. Received attribute values overwrite older ones.

&& The attribute's boolean value is determined by the incoming attribute values by using an *and*. For the aforementioned examples, the value is $i_1 \wedge i_2 \wedge i_3$.

|| The attribute's boolean value is determined by the incoming attribute values by using an *or*. For the aforementioned examples, the value is $i_1 \vee i_2 \vee i_3$.

***** The attribute's numeric value is determined by the incoming attribute values by multiplying them with each other. For the aforementioned examples, the value is $i_1 \cdot i_2 \cdot i_3$.

+ The attribute's numeric value is determined by the incoming attribute values by adding them to each other. For the aforementioned examples, the value is $i_1 + i_2 + i_3$.

Alias

The alias specifies how the corresponding attribute value is known in the simulator. For Rotalumis simulators, this requires both a process identifier and a name.

```
1 | in process1 as "difficultName"
```

For FMU simulators, it only requires a name.

```
1 | as "difficultName"
```

Example

A possible definition of a set of attributes is the following. Since the attributes *power* and *hasCollisions* do not have an alias, these attributes have identical names within the model that is used.

```
1 | Attributes {
2 |   InOutput Real power
3 |   Input Real actorXPosition as "x_position"
4 |   Output Boolean [Collision] hasCollisions
5 |   InOutput Boolean [||] activity in c as "activity"
6 | }
```

4.3.4 Parameters

Parameters are basically attributes that can be set during initialisation. These are defined by the following:

```
{DataType} {ID} "{Alias}" (in [Process])?
```

Again, the Process property should only be specified when the federate class is based on a POOSL model. The Alias property is a string that specifies the attribute name in the model. The DataType property is equal to the one in the Attribute definition as described above. Parameters are wrapped in a “Parameters” block. Samples are listed below.

```
1 | Parameters {
2 |   Real holdTime "HoldTime"
3 |   Real "HoldLevel"
4 |   Boolean active "Active"
5 |   Real VacantLevel "VacantLevel" in c
6 | }
```

4.3.5 Time policy

Sets the HLA time policy for this Federate Class. Value can be *RegulatedAndConstrained*, *Regulated*, *Constrained* or *None*. These values correspond with the HLA federate time policy, which can be found in the definition of the HLA standard [1].

```
1 | TimePolicy RegulatedAndConstrained
```

4.3.6 Default model

Where the federate class specification specifies the simulation model in terms of attributes and parameters, the default model to simulate may be specified as well. This can be done using the *DefaultModel* property. When specified, no further specification is required upon starting the co-simulation: the provided path to the FMU, POOSL model or collision model will be used as model to simulate. When it is used for BulletCollision simulator types, multiple models may be given by providing multiple file paths as strings. Two samples are the following.

```
1 | DefaultModel "../path/to/model/Model.fmu"
2 | DefaultModel "models/model1.json" "models/model2.json" "models/model3.json"
```

4.3.7 Advance type

AdvanceType is a property to provide a default step method to the federate. Options are the following. These properties also reflect different methods of advancing time according to the HLA standard [1].

- **TimeAdvanceRequest:** Request time advance grants to RTI.
- **NextMessageRequest:** Request next message requests to the RTI.

A sample definition is the following.

```
1 | AdvanceType NextMessageRequest
```

4.3.8 Default step size

`DefaultStepSize` can be used to specify the default step size for a simulator. The keyword is followed by a double value that indicates the step size. While using the `NextMessageRequest` method, this value represents the time-out to wait for a message. For POOSL models, this property does not have any effect, because these models provide their step size to the simulator. A sample is the following.

```
1 | DefaultStepSize 3.0
```

4.3.9 Default lookahead

The default lookahead value can also be specified using the `DefaultLookahead` property. The property is very similar to the step size and a sample is displayed below.

```
1 | DefaultLookahead 0.1
```

4.3.10 Simulation weight

The simulation weight represents the computational weight of the model simulation. A very computation intensive simulation should be assigned a higher weight than a very light simulation model. This value is used to create a distribution across multiple nodes automatically. The default value for each simulation is 1.

```
1 | SimulationWeight 3
```

4.4 ConnectionSets

`ConnectionSets` can be created for federate classes that typically have many instances being connected to each other by equal attribute connections. They specify the attribute connections on a class level instead of instance level as explained in Section 4.6.2, after which two instances can be connected to each other without having to specify every attribute connection between them again. A `ConnectionSet` starts with the `ConnectionSet` keyword, after which both federate classes are specified that use this `ConnectionSet` to be connected to each other. Then a block is started in which all attribute connections are specified. The layout of a `ConnectionSet` specification is as follows:

```
ConnectionSet between [FedClass] and [FedClass] { {Connection}+ }
```

4.4.1 ConnectionSet connection

Each of the `ConnectionSet`'s connections represents the connection on an attribute level: which attribute of one federate is connected to which attribute of the other federate. The format for this is the following:

```
{ [FedClass].[attr] <- [FedClass].[attr] }
```

Both attributes should be a member of the corresponding federate class. The first attribute receives its input from the second attribute. A sample `ConnectionSet` is displayed below.

```
1 | ConnectionSet between FedClassA and FedClassB {
2 |   { FedClassA.voltage <- FedClassB.outp }
3 |   { FedClassB.inp <- FedClassA.encoder }
4 | }
```

4.5 Configurations

A `Configuration` is a block definition containing a initialisation configuration for a federate class. It is defined like the following:

```
Configuration {ID} for [FederateClass] { {InitAssignment}+ }
```

The ID gives the configuration a name and the `FederateClass` specifies for which class the configuration is.

4.5.1 Initialisation attributes

Every `InitAssignment` assigns a value to a `Parameter` of the federate class. All values are put between quotes as string notation, regardless of the data type. An assignment looks like the following.

```
[Parameter] = "{Value}"
```

A sample `Configuration` for the parameters that were given in Section 4.3.4 is the following.

```
1 | Configuration Large for Room {
2 |     holdTime = "300.0"
3 |     holdLevel = "50.0"
4 |     active = "true"
5 |     VacantLevel = "10.0"
6 | }
```

4.6 Federations

Federations are used to specify a co-simulation. In the basis, it consists of a number of simulations (instances) and connections between them. A couple of extensions are added to provide some additional features. A `Federation` looks like the following:

```
Federation {ID} { {Instances}? {Connections}? ... }
```

4.6.1 Instances

A `Federation` consists of a number of simulations that are connected to each other. Each of these simulations is an `Instance` of a federate class as specified earlier. An instances block looks like the following:

```
Instances { ({ID} : [FederateClass])+ }
```

A sample instances block is the following.

```
1 | Instances {
2 |     sim1 : FedA
3 |     sim2 : FedA
4 |     sim3 : FedB
5 | }
```

4.6.2 Connections

Connections are similar to the `ConnectionSet` connections. Regular connections, however, connect attributes on an instance level. A connections block looks like the following:

```
Connections { {Connection}+ }
```

A connection can either be an attribute connection, an instance connection or a logger connection. An attribute connection connects two attributes from two instances to each other. Such a connection looks like the following:

```
{ [FedClass].[Attribute] <- [FedClass].[Attribute] }
```

An instance connection connects multiple attributes at once according to a previously defined `ConnectionSet`. The `ConnectionSet` to use does not need to be specified. An instance connection looks like the following:

```
{ [Instance] - [Instance] }
```

A logger connection connects multiple output attributes from different instances to a single logger instance. Such a connection looks like the following:

```
{ [Instance] <- [Instance].[Attribute] (, [Instance].[Attribute])* }
```

Altogether, a sample connections block could be the following.

```
1 | Connections {
2 |     { sim1.inp <- sim3.outp } // Attribute
3 |     { sim2.inp <- sim3.outp } // Attribute
4 |     { sim1 - sim2 } // Instance
5 |     { logger <- sim1.outp, sim2.outp } // Logger
6 | }
```


4.6.3 Situations

A **Situation** can be used to specify an initialisation configuration for a Federation. A Situation consists of a set of initialisation attribute assignments and/or applications of configurations to federate instances. Situations can extend each other to support several layers of initialisation configurations to be stacked. A Situation looks like the following:

```
Situation {ID} (extends [Situation])? { {SituationElement}+ }
```

Every SituationElement is either an attribute initialisation or an application of a model configuration.

- **Initialisation attribute initialiser**

The initialisation attribute initialiser specifies the value for one of the initialisation attributes for a specific federate instance in the Confederation. This SituationElement looks like the following:

```
Init [Instance].[InitAttribute] as "{Value}"
```

The Instance should already be defined in the Confederation and only initialisation attributes can be assigned. The value must be specified as a String.

- **Application of model configuration**

An application of a model configuration applies a model configuration to a specific federate instance. This is particularly useful when assigning multiple initialisation attributes at once. This SituationElement looks like the following:

```
Apply [ModelConfiguration] to [Instance]
```

An extending Situation might overwrite initialisation attribute values of the Situation it extends. A sample Situations block might look like the following.

```
1 | Situation sit1 {
2 |   Apply large to room1
3 |   Init room1.heaterSize as "3.2"
4 |   Apply small to room2
5 |   Init room2.heaterSize as "1.1"
6 | }
7 | Situation sit2 extends sit1 {
8 |   Apply medium to room1
9 |   Init room1.heaterSize as "2.3"
10| }
```

4.6.4 Scenarios

A **Scenario** is a predefined sequence of events taking place during the co-simulation. A scenario optionally starts with an end time specified for when the simulation should be stopped:

AutoStop: ({FLOAT} | no) Every **Scenario** has an ID and supports both assignments of attributes in the co-simulation and messages being sent over sockets. For the latter case, sockets should be specified first in the **Scenario** block. The targeted instance connects to the specified socket after starting the co-simulation. A socket specification looks like the following:

```
Socket {ID} for [Instance] on "{Hostname}":{Port}
```

After the (optional) specification of the sockets, the sequence of events is specified. Every **Event** occurs at a given time and is either an assignment or a socket message (in bytes). An assign event looks like the following:

```
On {Time} Assign "{Value}" to [Instance].[Attribute]
```

A socket event looks like the following:

```
On {Time} Send [Socket] bytes {Byte}+ (, {Byte}+)?
```

Here, the Socket refers to one of the previously defined sockets and Byte represents a sequence of bytes in hexadecimal representation, such as 0x01 0xe7 0xB1. A sample Scenarios block is the following.

```
1 | Scenario sampleScenario {
2 |   AutoStop: 30.0
3 |   Socket sock1 for room1 on "localhost":10001
4 |   Socket sock2 for room2 on "localhost":10002
5 |   On 10.0 Send sock1 bytes 0x01 0x02 0x03
6 |   On 10.0 Send sock2 bytes x01 x02 x03           // Byte notation is allowed too
```

```

7 | On 15.0 Assign "false" to room1.heaterState
8 | }

```

4.6.5 Fault scenarios

A federation might contain a number of `FaultScenario`s. Every `FaultScenario` looks like the following:

```
FaultScenario { {Fault}+ }
```

For most of the faults, either a single moment in time, the period starting from a specified time, or a time range can be specified. This definition is called a **Range** and looks like the following:

```
(On {Time}) | (From {Time} (to {Time})?)
```

There are different types of Faults that could occur; these are listed below.

- **AccuracyFault:** An `AccuracyFault` introduces some variance on an attribute. It looks like the following:

```
Variance for [Instance].[Attribute] = {Float}
```
- **TimedAbsoluteFault:** A `TimedAbsoluteFault` sets the value of an attribute to a predefined value at a specific time. It looks like the following:

```
{Range} set [Instance].[Attribute] = "{Value}"
```
- **TimedConnectionFault:** A `TimedConnectionFault` interrupts the connection of an attribute to the other attributes, disabling it to publish its most recent values. It looks like the following:

```
{Range} disconnect [Instance].[Attribute]
```
- **TimedOffsetFault:** A `TimedOffsetFault` introduces an offset to an attribute value; it looks like the following:

```
{Range} offset [Instance].[Attribute] = "{Value}"
```

All `Value` elements in parenthesis are String representations of the attribute value to be assigned.

```

1 | FaultScenario sampleFaultScenario {
2 |   Variance for fed1.attr2 = 0.3
3 |   From 2.0 to 5.0 set fed2.attr1 = "2.5"
4 |   From 12.0 offset fed4.attr3 = +0.2
5 |   On 8.0 disconnect fed3.attr1
6 | }

```

4.6.6 Design Space Exploration

Every Design Space Exploration (DSE) has an ID and some basic settings: `SweepMode`, `Scenario` and `FaultScenario`. These are all optional. The `Scenario` and `FaultScenario` properties specify which of these settings to run for each of the design space configurations when co-simulating. The `SweepMode` specifies how the configurations should be swept, either `Independent` (default) or `Linked`. Using independent sweep mode, all parameter values are used with all other parameter values, while linked mode only runs the n-th value of a parameter with all other n-th values of the other parameters. Consequently, a DSE of only two parameters, each having three possible values, will result in 9 runs using independent sweep mode and only 3 runs using linked sweep mode.

Hereafter, the design space is specified by providing `Situations`, `Configurations` and `Initialisation Attributes` to sweep. Each of these lists is specified as listed below.

- `Situations:` `[Situation] (, [Situation])*`
- `Configurations for [Instance] :` `[Configuration] (, [Configuration])*`
- `Set [Instance].[InitAttribute]:` `{Value} (, {Value})*`

A Sample DSEs block looks like the following.

```

1 | DSE dse1 {
2 |   SweepMode Independent
3 |   Scenario sampleScenario
4 |   Faults sampleFaultScenario
5 |   Situations: sit1, sit2
6 |   Configurations for room1: small, medium, large
7 |   Configurations for room2: small, medium, large
8 |   Set room1.heaterSize: 1.2, 2.3, 3.4
9 |   Set room2.heaterSize: 1.1, 2.2
10| }

```

The sample DSE “dse1” has a design space of 108 configurations.

4.6.7 Metric sets

A federation may consist zero or more `MetricSets`. Every `MetricSet` has an ID and requires a `MeasureTime` to be defined. This looks like the following:

```
MetricSet {ID} { MeasureTime: {Float} {Metric}+ }
```

Every `Metric` also has an ID and can either be one of the following metrics.

- **EndValue**

The `EndValue` metric returns the last value of the specified attribute. This value can optionally be the absolute value and may be relative to another attribute. In the latter case, the difference between the attributes is calculated and returned.

```
Absolute? EndValue [Instance].[Attribute] (relative to [Instance].[Attribute])?
```

- **Error**

The `Error` metric calculates the mean error between two attributes and returns this error. Optionally, this may be set to be a squared error value.

```
Squared? Error [Instance].[Attribute] (relative to [Instance].[Attribute])?
```

- **Timer**

The `Timer` metric returns the time which is the first moment during the simulation for which the provided expression holds. Such a metric could, for example, be used to find the first moment for which a temperature exceeds a specific value. A timer metric may be specified to be an end condition for the collection of metrics. When all metrics that have been set to be end conditions have finished (have a first moment in time for which the condition holds), the co-simulation is stopped and all other metric values are calculated and returned, together with the results of the timer metrics.

```
Timer for [Instance].[Attribute] {Comparator} {Value} (EndCondition (after {Time})?)?
```

- **Minimum**

The `minimum` metric returns the minimum value reached by the specified attribute during the simulation.

```
Minimum of [Instance].[Attribute]
```

- **Maximum**

The `maximum` metric returns the maximum value reached by the specified attribute during the simulation.

```
Maximum of [Instance].[Attribute]
```

`Absolute` and `Squared` are optional extensions to the `EndValue` and `Error` metrics respectively, as well as the `relative to` property, which allows the user to request a metric value relative to another attribute value. The `Comparator` can be one of the following: `<`, `<=`, `==`, `>=`, `>`, `!=`. Once the condition made by comparing the attribute value to the specified value using the comparator yields `true`, the time is set a metric result for the timer. Optionally, the simulation may be triggered to exit when this timer yields `true`. A delay before closing can also be specified. `Minimum` and `maximum` value are self-explaining.

4.6.8 Distributions

Distribution configurations for use of distributed simulation over multiple nodes may be specified for a federation. Every `Distribution` has an ID and a number of nodes specified. This looks like the following:

```
Distribution {ID} over {Number} systems { {SystemSet}+ }
```

Every `SystemSet` is assigned an identification number of the node and a list of federate instances that should be executed on the node. This looks like the following:

```
System {Number}: [Instance]+
```

A sample of a `DistributionSets` block is the following.

```
1 | Distribution dist1 over 2 systems {
2 |     System 0: fed1 fed2 fed3
3 |     System 1: fed4 fed5 fed6
4 | }
5 | Distribution dist2 over 3 systems {
6 |     System 0: fed1 fed4
7 |     System 1: fed2 fed3 fed5
8 |     System 2: fed6
9 | }
```

5 Running a co-simulation

Once a Confederation has been defined, a CMake project, configuration files and run-script are generated in the Source Generation Directory.

5.1 Configuration files

The `conf` directory contains all configuration files divided in a couple of subdirectories. Each of these directories and their contents are listed below. For each `FederateClass` having Model Configurations specified, a new directory is created.

- `conf`
 - Topology file holding information on instances and connections (`.topo`)
- `conf/[Federation]`
 - Situation configurations (`.situation`)
 - Scenarios (`.scenario`)
 - Fault scenarios (`.fscenario`)
 - Design Space Exploration configurations (`.dse`)
 - Metric sets (`.metrics`)
 - Distributions (`.distribution`)
- `conf/[FederateClass]`
 - Model configurations (`.init`)

5.2 Run-script

In the root of the Source Generation Directory, a Python run-script (`run.py`) is generated. This run-script can be used to build and run the co-simulation. The script must be executed from the command line and supports a set of flags and properties to specify its actions. All options are listed below and are categorised. A number of useful commands is provided in Section 5.2.6.

5.2.1 Basic behaviour

These flags tell the script what to do.

- **-b** Build the generated sources using CMake.
- **-d** Display the active configuration.
- **-e** Start the simulation using the active configuration.
- **-h** Display the help message.

5.2.2 Configurables

These properties can be used to configure the co-simulation.

- **-t *x*, --topology *x***
Sets the topology configuration file to *x*.
- **-f *x*, --fom *x***
Sets the FOM XML file to *x*.
- **-r *x*, --rti *x***
Sets the RTI address (hostname) to *x*.
- **-s *x*, --situation *x***
Sets the situation configuration file to *x*.
- **-u *x*, --scenario *x***
Sets the user scenario file to *x*.
- **-m *x*, --metrics *x***
Sets the configuration files for metrics to *x*.
- **--faults *x***
Sets the fault scenario file to *x*.
- **--dse *x***
Sets the Design Space Exploration configuration file to *x*.

5.2.3 Distribution setup

These options can be used to start a distributed co-simulation.

- **--distribution *x***
Set the predefined distribution configuration file to *x*.
- **--systemId *x***
Set the node/system ID of the current node to *x*.
- **--nrOfSystems *x***
Set the total number of nodes in the distribution execution to *x*.
- **--parentHost *x***
Set the parent hostname to *x*.
- **--parentPort *x***
Set the parent port to *x*.

5.2.4 Other options

A couple of other options to apply changes are available.

- `-v, --verbose`
If set, all outputs from external commands – such as CMake – will be displayed in the console.
- `--disable x`
Disable the execution of federate *x*.
- `--cmake=x`
Add additional flags *x* to the CMake build command.
- `--render`
If set, collision simulators will also render the provided models.

5.2.5 Federate-specific options

Since all federates also have their own options and parameters, the exact set of command line arguments depends on the topology that is used. Arguments such as the measure time, step size, lookahead and models to use for the co-simulation depend on the used instances. Also, parameter values can be set by providing values to the run-script. To retrieve a list of all possible arguments, a topology should be provided to the run-script together with the `-h` option to print the help message for the specific topology. An example to request all options from the run-script is the following.

```
1 | ./run.py -t conf/sampleFederation.topo -h
```

5.2.6 Useful commands

A number of useful commands are listed below.

```
1 | // Display help message
2 | ./run.py -h
3 |
4 | // Display help message for a specific topology, including all arguments and options
5 | ./run.py -t conf/sampleFederation.topo -h
6 |
7 | // Build all sources and display output
8 | ./run.py -bv
9 |
10 | // Build and start the basic configuration
11 | ./run.py -t conf/sampleFederation.topo -be
12 |
13 | // Start a specific configuration of the system
14 | ./run.py -t conf/sampleFederation.topo -s conf/SampleFederation/configuration.
    | situation -e
15 |
16 | // Set situation and fault scenario and start the co-simulation
17 | ./run.py -t conf/sampleFederation.topo -u conf/SampleFederation/WorkDay.scenario -s
    | conf/SampleFederation/Medium.situation --faults conf/SampleFederation/errors.
    | fscenario -e
```

References

- [1] IEEE Standard for Modeling and Simulation (M&S) High Level Architecture (HLA)– Framework and Rules. *IEEE Std 1516-2010*, pages 1–38, Aug 2010.

A CoHLA grammar

```
1  /*
2  * Copyright (c) Thomas Nägele and contributors. All rights reserved.
3  * Licensed under the MIT license. See LICENSE file in the project root for details.
4  */
5
6  grammar nl.ru.sws.cohla.CoHLA with org.eclipse.xtext.common.Terminals
7
8  import "http://www.eclipse.org/emf/2002/Ecore" as ecore
9
10 generate coHLA "http://www.ru.nl/sws/cohla/CoHLA"
11
12 Model:
13   (
14     imports+=Import*
15     environment=Environment?
16     federateObjects+=FederateObject*
17     interfaces+=Interface*
18     configurations+=ModelConfiguration*
19     federations+=Federation*
20   );
21
22 Import:
23   'import' importURI=STRING
24   ;
25
26 // ===== HLA Environment =====
27
28 Environment:
29   'Environment' '{'
30     rti=HLAImplementation
31     ('SourceDirectory' srcDir=STRING)?
32     ('PrintLevel' printLevel=PrintLevel)?
33     ('publishOnlyChanges' ?= 'PublishOnlyChanges')?
34   '}'
35   ;
36
37 HLAImplementation:
38   'RTI' '{'
39     implementation=HLAImp
40     'Libraries' libRoot=STRING
41     ('Dependencies' depRoot=STRING)?
42   '}'
43   ;
44
45 enum HLAImp:
46   pitchRti="PitchRTI"
47   | portico="Portico"
48   | openRti="OpenRTI"
49   ;
50
51 enum PrintLevel:
52   state="State"
53   | time="Time"
54   | none="None"
55   ;
56
57 // ===== FederateObject =====
58
59 FederateObject:
60   'FederateClass' name=ID '{'
61     'Type' type=FederateType
62     ('{ config=SimulatorConfig }')?
63     ('Attributes' '{ attributes+=Attribute+ }')?
64     ('Parameters' '{ parameters+=Parameter+ }')?
65     ('TimePolicy' timePolicy=TimePolicy)?
66     ('DefaultModel' defaultModel+=STRING+)?
67     ('AdvanceType' advanceType=AdvanceType)?
68     ('DefaultStepSize' defaultStepSize=P_FLOAT)?
```

```

69     ('DefaultLookahead' defaultLookahead=P_FLOAT)?
70     ('SimulationWeight' simulationWeight=INT)?
71     '}',
72 ;
73
74 enum TimePolicy:
75     both="RegulatedAndConstrained"
76     | regulated="Regulated"
77     | constrained="Constrained"
78     | none="None"
79 ;
80
81 Process:
82     name=ID 'in' path=STRING
83 ;
84
85 Parameter:
86     dataType=DataType name=ID alias=STRING ('in' process=[Process])?
87 ;
88
89 SimulatorConfig:
90     POOSLConfig | LoggerConfig
91 ;
92
93 POOSLConfig:
94     'Processes' '{'
95     processes+=Process+
96     '}'
97 ;
98
99 LoggerConfig:
100     'DefaultMeasureTime' defaultMeasureTime=P_FLOAT
101 ;
102
103 enum FederateType:
104     poosl="POOSL"
105     | fmu="FMU"
106     | csv="CSV-logger"
107     | colsim="BulletCollision"
108     | none="None"
109 ;
110
111 enum AdvanceType:
112     time="TimeAdvanceRequest"
113     | message="NextMessageRequest"
114 ;
115
116 // ===== Attribute =====
117
118 Attribute:
119     sharingType=SharingType dataType=DataType (collision?='[Collision]')? ('['
120         multiInputOperator=MultiInputOperator ']'?)? name=ID (alias=AliasProperty |
121         processProperty=ProcessProperty)? ('{'
122         updateTypeProperty=UpdateTypeProperty?
123         updateConditionProperty=UpdateConditionProperty?
124         transportationProperty=TransportationProperty?
125         orderProperty=OrderProperty?
126         '}'?)?
127 ;
128
129 enum SharingType:
130     neither="Void"
131     | publish="Output"
132     | subscribe="Input"
133     | publishSubscribe="InOutput"
134 ;
135
136 enum DataType:
137     integer="Integer"
138     | long="Long"

```



```

137 | string="String"
138 | real="Real"
139 | bool="Boolean"
140 ;
141
142 ProcessProperty:
143   'in' process=[Process] 'as' attributeName=STRING
144 ;
145
146 AliasProperty:
147   'as' alias=STRING
148 ;
149
150 UpdateTypeProperty:
151   'UpdateTypeProperty' updateType=UpdateType
152 ;
153
154 UpdateConditionProperty:
155   'UpdateCondition' updateCondition=UpdateCondition
156 ;
157
158 enum UpdateCondition:
159   onChange="OnChange"
160 ;
161
162 TransportationProperty:
163   'Transportation' transportation=TransportationType
164 ;
165
166 OrderProperty:
167   'Order' order=OrderType
168 ;
169
170 // ===== ConnectionSet =====
171
172 Interface:
173   'ConnectionSet' 'between' class1=[FederateObject] 'and' class2=[FederateObject] '{'
174     connections+=InterfaceConnection+
175   '}'
176 ;
177
178 InterfaceConnection:
179   '{'
180     in=ClassAttributeReference '<-' out=ClassAttributeReference
181   '}'
182 ;
183
184 ClassAttributeReference:
185   objectClass=[FederateObject] '.' attribute=[Attribute]
186 ;
187
188 // ===== Federation =====
189
190 Federation:
191   'Federation' name=ID '{'
192     ('Instances' '{'
193       instances+=FederateInstance+
194     '}')?
195     ('Connections' '{'
196       connections+=Connection+
197     '}')?
198     situations+=Situation*
199     scenarios+=Scenario*
200     faultScenarios+=FaultScenario*
201     dses+=DSEConfig*
202     metricSets+=MetricSet*
203     distributions+=Distribution*
204   '}'
205 ;
206

```

```

207 FederateInstance:
208     name=ID ':' federate=[FederateObject]
209 ;
210
211 Connection:
212     AttributeConnection
213     | LoggerConnection
214     | ObjectConnection
215 ;
216
217 AttributeConnection:
218     '{'
219     inAttribute=AttributeReference '<-' outAttributes+=AttributeReference
220     '}'
221 ;
222
223 LoggerConnection:
224     '{'
225     logger=[FederateInstance] '<-' outAttributes+=AttributeReference (','
226     outAttributes+=AttributeReference)*
227     '}'
228 ;
229
230 ObjectConnection:
231     '{'
232     instance1=[FederateInstance] '->' instance2=[FederateInstance]
233     '}'
234 ;
235
236 AttributeReference:
237     instance=[FederateInstance] '.' attribute=[Attribute]
238 ;
239 // ===== Situation =====
240
241 Situation:
242     'Situation' name=ID ('extends' extends=[Situation])? '{'
243     elements+=SituationElement+
244     '}'
245 ;
246
247 SituationElement:
248     AttributeInitialiser | ApplyModelConfiguration
249 ;
250
251 AttributeInitialiser:
252     'Init' reference=InstanceParameterReference 'as' value=STRING
253 ;
254
255 InstanceParameterReference:
256     instance=[FederateInstance] '.' attribute=[Parameter]
257 ;
258
259 ApplyModelConfiguration:
260     'Apply' configuration=[ModelConfiguration] 'to' instance=[FederateInstance]
261 ;
262
263 // ===== Configuration =====
264
265 ModelConfiguration:
266     'Configuration' name=ID 'for' federateObject=[FederateObject] '{'
267     initAttributes+=AttributeConfigurator+
268     '}'
269 ;
270
271 AttributeConfigurator:
272     reference=ParameterReference '=' value=STRING
273 ;
274
275 ParameterReference:

```

```

276     attribute=[Parameter]
277 ;
278
279 // ===== Scenarios =====
280
281 Scenario:
282     'Scenario' name=ID '{'
283         settings=ScenarioSettings?
284         events+=Event+
285     '}'
286 ;
287
288 ScenarioSettings:
289     'AutoStop:' (autoStopTime = P_FLOAT | noAutoStop ?= 'no')
290     sockets+=ScenarioSocket*
291 ;
292
293 ScenarioSocket:
294     'Socket' name=ID 'for' instance=[FederateInstance] 'on' host=STRING ':' port=INT
295 ;
296
297 Event:
298     time=P_FLOAT ':' action=Action
299 ;
300
301 Action:
302     ActionAssign | ActionSocket
303 ;
304
305 ActionAssign:
306     attribute=AttributeReference '=' value=STRING
307 ;
308
309 ActionSocket:
310     socket=[ScenarioSocket] '<' bytes+=HexByte (',' bytes+=HexByte)*
311 ;
312
313 // ===== Faults =====
314
315 FaultScenario:
316     'FaultScenario' name=ID '{'
317         faults+=Fault+
318     '}'
319 ;
320
321 Fault:
322     TimedAbsoluteFault | TimedConnectionFault | TimedOffsetFault | AccuracyFault
323 ;
324
325 AccuracyFault:
326     'Variance' 'for' attribute=AttributeReference '=' value=P_FLOAT
327 ;
328
329 TimedAbsoluteFault:
330     range=Range 'set' attribute=AttributeReference '=' value=STRING
331 ;
332
333 TimedConnectionFault:
334     range=Range 'disconnect' attribute=AttributeReference
335 ;
336
337 TimedOffsetFault:
338     range=Range 'offset' attribute=AttributeReference '=' value=C_FLOAT
339 ;
340
341 Range:
342     ('On' time=P_FLOAT)
343     | ('From' startTime=P_FLOAT ('to' endTime=P_FLOAT)?)
344 ;
345

```

```

346 // ===== Domain Space Exploration =====
347
348 DSEConfig:
349   'DSE' name=ID '{'
350     ('SweepMode' sweepMode=SweepMode)?
351     ('Scenario' scenario=[Scenario])?
352     ('Faults' faultScenario=[FaultScenario])?
353     situations+=DSESituationConfig*
354     configurations+=DSEFederateConfig*
355     attributes+=DSEAttrInit*
356   '}'
357 ;
358
359 DSESituationConfig:
360   'Situations' ':' situations+=[Situation] (',' situations+=[Situation])*
361 ;
362
363 DSEFederateConfig:
364   'Configurations' 'for' instance=[FederateInstance] ':' configurations+=[
365     ModelConfiguration] (',' configurations+=[ModelConfiguration])*
366 ;
367
368 DSEAttrInit:
369   'Set' attr=InstanceParameterReference ':' value=MultiValue
370 ;
371
372 enum SweepMode:
373   independent="Independent"
374   | linked="Linked"
375 ;
376
377 MultiValue:
378   value=(P_FLOAT | C_FLOAT | C_INT | STRING) | values+= (P_FLOAT | C_FLOAT | C_INT |
379     STRING) (',' values+=(P_FLOAT | C_FLOAT | C_INT | STRING))+
380 ;
381
382 // ===== Metrics =====
383
384 MetricSet:
385   'MetricSet' name=ID '{'
386     'MeasureTime:' measureTime=P_FLOAT
387     metrics+=Metric+
388   '}'
389 ;
390
391 Metric:
392   'Metric' name=ID 'as' (metricEV=MetricEndValue | metricErr=MetricError |
393     metricTimer=MetricTimer | metricMinMax=MetricMinMax)
394 ;
395
396 MetricEndValue:
397   absolute?='Absolute'? 'EndValue' attribute=AttributeReference relativeTo=
398     MetricRelativeTo?
399 ;
400
401 MetricError:
402   squared?='Squared'? 'Error' attribute=AttributeReference relativeTo=
403     MetricRelativeTo
404 ;
405
406 MetricTimer:
407   'Timer' 'for' attribute=AttributeReference comparator=Comparator value=(P_FLOAT |
408     C_FLOAT | C_INT | 'true' | 'false') ((' isEndCondition?='EndCondition' ('after
409     ' delay=P_FLOAT)? '))'?
410 ;
411
412 MetricMinMax:
413   (min?='Minimum' | max?='Maximum') 'of' attribute=AttributeReference
414 ;
415

```

```

409 MetricRelativeTo:
410   'relative' 'to' ref=AttributeReference
411 ;
412
413 // ===== Distributions =====
414
415 Distribution:
416   'Distribution' name=ID 'over' nrOfSystems=INT 'systems' '{'
417     systemSets+=SystemSet+
418   '}'
419 ;
420
421 SystemSet:
422   systemId=INT ':' federates+=[FederateInstance]+
423 ;
424
425 // ===== Types =====
426
427 HexByte:
428   bytes+=HEXPAIR+
429 ;
430
431 enum OrderType:
432   receive="Receive"
433   | timeStamp="TimeStamp"
434 ;
435
436 enum TransportationType:
437   reliable="Reliable"
438   | bestEffort="BestEffort"
439 ;
440
441 enum UpdateType:
442   conditional="Conditional"
443 ;
444
445 enum Comparator:
446   lt="<"
447   | le="<="
448   | eq="=="
449   | ge=">="
450   | gt=">"
451   | ne="!="
452 ;
453
454 enum MultiInputOperator:
455   none = "NONE"
456   | and = "&&"
457   | or = "||"
458   | plus = "+"
459   | product = "*"
460 ;
461
462 terminal HEXPAIR returns ecore::EString:
463   ('0x'|'x') HEXCHAR HEXCHAR
464 ;
465
466 terminal HEXCHAR returns ecore::EChar:
467   ('0'..'9'|'A'..'F'|'a'..'f')
468 ;
469
470 terminal P_FLOAT returns ecore::EString:
471   INT '.' ('0'..'9')+
472 ;
473
474 terminal C_FLOAT returns ecore::EString:
475   C_INT '.' ('0'..'9')+
476 ;
477
478 terminal C_INT returns ecore::EString:

```

```
479 | ('-' | '+') INT  
480 | ;
```