ResearchGate

See discussions, stats, and author profiles for this publication at: http://www.researchgate.net/publication/286455765

Easing the development of HLA Federates: the HLA Development Kit and its exploitation in the SEE Project.

CONFERENCE PAPER · OCTOBER 2015

READS
12

6 AUTHORS, INCLUDING:

Alberto Falcone
Università della Calabria

7 PUBLICATIONS **3** CITATIONS

SEE PROFILE

Alfredo Garro
Università della Calabria

61 PUBLICATIONS **468** CITATIONS

SEE PROFILE

Nauman Riaz Chaudhry
University of Gujrat

3 PUBLICATIONS **1** CITATION

SEE PROFILE

All in-text references underlined in blue are linked to publications on ResearchGate, letting you access and read them immediately.

Available from: Alberto Falcone
Retrieved on: 22 December 2015

Easing the development of HLA Federates: the HLA Development Kit and its exploitation in the SEE Project

Alberto Falcone, Alfredo Garro
Department of Informatics, Modeling, Electronics and Systems Engineering (DIMES)
University of Calabria
Rende (CS), Italy
{firstname.lastname}@dimes.unical.it

Anastasia Anagnostou, Nauman R. Chaudhry, Omar-Alfred Salah, Simon J.E. Taylor
Department of Computer Science
University of Brunel
Uxbridge, United Kingdom
{firstname.lastname}@brunel.ac.uk

***Abstract*— The Modeling & Simulation (M&S) of modern cyber-physical systems is presenting new challenges. New M&S techniques, methods and tools are emerging that take advantage of distributed simulation environments. One of the most mature and popular standard for distributed simulation is the IEEE 1516-2010 - High Level Architecture (HLA) that, although originally developed for military applications, is increasingly exploited in a great variety of application domains due to its capabilities to enable the interoperability and reusability of distributed simulation components. However, the development of fully fledged simulation models, based on the IEEE 1516-2010 standard, is still a challenging task and requires considerable development effort that often results not only in an increase in development time but also in low reliability. In this context, the paper presents a general-purpose, domain-independent framework that aims to ease the development of HLA-based simulations. Its effectiveness is exemplified in the context of the Simulation Exploration Experience (SEE) project lead by NASA and which involves several U.S. and European Institutions.**

Keywords*— *Distributed Simulation; High Level Architecture; Agent-based Simulation.

I. Introduction

Modeling and Simulation (M&S) represents one of the most important and effective methods for designing and studying complex systems in a variety of industrial and scientific domains ranging from biology to space exploration [10]. M&S methods, tools, and techniques allow effectively analyzing and evaluating design alternatives by avoiding risks, costs and fails associated with extensive field experimentation; this opportunity becomes crucial when complete and actual tests are too expensive to be performed in terms of cost, time and other primary resources.

Over the years, large-scale systems have increased in complexity and sophistication since, in general, they are composed of several components, which are often designed and developed by organizations belonging to different engineering domains, including mechanical, electrical, and software. As systems get increasingly complex, their design and development become more difficult and therefore new M&S techniques, methods and tools are emerging also to benefit from distributed simulation environments [1], [4]. In this context, the IEEE 1516-2010 - High Level Architecture (HLA) standard [5] attempts to handle this complexity by providing a specification of a distributed infrastructure in which simulation units can run on standalone computers and communicate with one another in a common simulation scenario.

HLA development was initiated and sponsored by the U.S. Department of Defense to facilitate the integration of distributed simulation models within a common architecture. Although it was initially developed to support military applications, it has been widely used in non-military industries for its many advantages related to the interoperability and reusability of distributed simulation components. In the HLA standard a distributed simulation is called a *Federation,* and it is composed of several HLA simulation entities, each called a *Federate*, which can interact among them by using a Run-Time Infrastructure (RTI). The RTI represents a backbone of a Federation execution that provides a set of standard protocols and services to manage the communications and data exchange among Federates. Each Federation has a Federation Object Model (FOM) that is created in accordance with the Object Model Template (OMT) defined by the standard [5].

Building complex and large distributed simulations systems, based on the IEEE 1516-2010 standard, is a challenging task and requires considerable development efforts. Indeed, it requires expert engineers with knowledge and experience in distributed systems, simulation, middleware and software programming. The main problem is that the development and testing of HLA Federates are generally difficult, complex, and resource-intensive not only because of the complexity of the IEEE 1516-2010 family standards [5] but also due to the lack of proper documentations and ready-to-use examples. Moreover, developers have to spend a considerable effort to handle common HLA functionalities, such as the management of the simulation time, the connection on the RTI, and the management of common RTI exceptions. As a result, they cannot fully focus on the specific aspects of their own simulations (the HLA Federates). Thus, it would be desirable to separate the common HLA aspects from those of a specific

HLA Federate by providing a general-purpose, domain-independent framework that allows achieving these goals.

In this context, the paper presents the HLA Development Kit, a general-purpose and domain-independent toolkit that eases the development of HLA Federates by providing a software framework, called DKF (Development Kit software Framework), with related documentation, user guide and reference examples. Specifically, the DKF allows developers to focus on the specific aspects of their own HLA Federates rather than dealing with the common HLA functionalities, which are managed by the DKF core components.

The rest of the paper is organized as follows. Related works are discussed in Section II. Section III presents the HLA Development Kit with particular focus on the architecture and main services provided by the HLA Development Kit software Framework (DKF). In Section IV, the development of a HLA Federate from scratch based on the DKF is exemplified in the context of the Simulation Exploration Experience (SEE) project. Qualitative and quantitative analysis of the benefits provided by the DKF are presented in Section V and VI respectively. Finally, conclusions are drawn and future research directions are presented.

II. RELATED WORK

Several research efforts focused their attention on the creation of HLA simulation and development environments, mainly aiming at providing an integrated toolchain for creating and simulating complex systems by using specialized modeling tools and methodologies.

For MATLAB/Simulink different packages and toolboxes are available for implementing HLA simulators such as the *Forwardsim HLA Toolbox for MATLAB* [13], which provides a user interface that allows developers to fully design and customize their HLA Federates. Another tool is the *HLA/DIS Toolbox for MATLAB and Simulink* [6] that provides a library of Simulink blocks specifically designed for the integration of High Level Architecture services into Simulink models. It greatly simplifies Federation development and model reuse, as well as enables organizations to more efficiently participate in multinational simulations or implement distributed simulation models locally.

Another tool that enables developers to effectively manage the structure and assets of a HLA Federate starting from a FOM file is the *PITCH Developer Studio* [7]. This software allows programmers to reduce the HLA learning curve by providing functionalities for creating and exporting auto-generate C++/Java code classes based on the structure of the HLA Federate.

A domain-specific HLA software framework was created by the Danish Maritime Institute (DMI) [14]. This framework defines a universe of real-time simulation concepts to support the more informal concepts available at DMI with a HLA environment. The simulation framework provides mechanisms to simplify the development of real-time simulators.

Other HLA frameworks are based on GRID-computing infrastructure and they have become in recent times a popular way to model and study complex multi-actor systems by using the typical characteristics and capabilities of the GRID [15].

The HLA Development Kit and its software framework (DKF) presented in this paper differ from the above mentioned solutions in several aspects. In particular, differently from a proprietary and commercial solution that requires tool-specific knowledge and training, the HLA Development Kit is an open source project released under the open source policy Lesser GNU Public License (LGPL) and can be freely and easily customized and/or extended to cover and deal with both domain independent and domain-specific aspects (as was the case with the SEE-specific extension presented in Section IV). In addition, the DKF provides advanced facilities that allow keeping the code compact, readable and reliable (see Section VI). As an example, Java annotations are used to directly inject the structure of a HLA Federate in the Java code. These metadata are used by the core components of the DKF at run-time to inspect and check a HLA Federate according to its definition in the FOM. The above sketched capabilities showed their great benefits not only for expert HLA developers but also for HLA novice practitioners as were the undergraduates students involved in the SEE project (see Section V).

III. THE HLA DEVELOPMENT KIT

As sketched in Section I, the development of a Federate and a Federation is complex and there are few tutorial resources to help educate developers. In order to help new developers, the *HLA Development Kit*, which provides high level functionality both to implement HLA Federates and manage the interactions between them and the RTI, has been developed.

The *HLA Development Kit* aims at easing the development of HLA Federates by providing the following resources: (i) a software framework (the DKF) for the development in Java of HLA Federates; (ii) a technical documentation that describes the DKF; (iii) a user guide to support developers in the use of the DKF; (iv) a set of reference examples of HLA Federates created by using the DKF; and, (v) video-tutorials, which show how to create both the structure and the behavior of a HLA Federate by using the DKF.

In the following, the attention is focused on the DKF and, specifically, on its architecture and underlying Federate model-behavior. Moreover, a domain-specific extension of the DKF is also presented.

A. The HLA Development Kit Framework (DKF)

The DKF is a general-purpose, domain-independent framework, released under the open source policy Lesser GNU Public License (LGPL), which facilitates the development of HLA Federates [5]. Indeed, the DKF allows developers to focus on the specific aspects of their own Federates rather than dealing with the common HLA functionalities, such as the management of the simulation time; the connection/disconnection on/from the HLA RTI; the publishing, subscribing and updating of *ObjectClass* and *InteractionClass* elements. The DKF is designed and developed by the SMASH-Lab (System Modeling And Simulation Hub - Laboratory) of the University of Calabria (Italy) working in

cooperation with the NASA JSC (Johnson Space Center), Houston (TX, USA).

The DKF is fully implemented in the Java language and is based on the following three principles: (i) *Interoperability*, DKF is fully compliant with the IEEE 1516-2010 specifications; as a consequence, it is platform-independent and can interoperate with different HLA RTI implementations (e.g. PITCH, VT/MÄK, PoRTIco, CERTI); (ii) *Portability and Uniformity*, DKF provides a homogeneous set of APIs that are independent from the underlying HLA RTI and Java version. In this way, developers could decide the HLA RTI and the Java run-time environment at development-time; and (iii) *Usability*, the complexity of the features provided by the DKF framework are hidden behind an intuitive set of APIs.

The design and implementation of the DKF has been centered on typical Software Engineering methods and, in particular, on an *agile* software development process. Furthermore, it has been developed according to the concept of *Object HLA*, in this way, the development of HLA Federates could benefit also from the *Object HLA* features and functionalities provided by the Pitch Developer Studio [7] or similar IDE.

To promote the adoption and experimentation of the HLA Development Kit and its DKF, the Kit has been specialized in the *SEE HLA Starter Kit* with the aim to ease the development of HLA Federates in the context of the Simulation Exploration Experience (SEE) project [9]. SEE is an event organized by the Simulation Interoperability Standards Organization (SISO), in collaboration with NASA and other research and industrial partners, with the objective to promote the adoption of the HLA standard and compliant tools by involving university teams in the distributed simulation of a Moon base. The SEE-specific features introduced in both the DKF and the Development Kit (as an example, the implementation of SEE Dummy and Tester Federates) aim not only at reducing the development efforts but also at improving the reliability of SEE Federates and thus reducing the problems arising during the final integration and testing phases of the SEE project [9]. Moreover, this SEE extension allows to prove how, starting from a domain-independent core of the DKF, conceived for supporting the development of general-purpose HLA Federate, it is possible to easily add and integrate application-specific extensions for supporting the development of domain-specific Federates.

The following subsections are devoted to present both the architectural and behavioral aspects of the DKF also with reference to its SEE-specific extension (the SEE-SKF: *SEE HLA Starter Kit Framework*).

1) Architecture of the DKF

The architecture of a DKF-based Federation is composed of three main layers (see Fig. 1): (i) *Application Layer*, which contains the Federates that can interact with both the DKF and the HLA RTI by using their APIs; (ii) *DKF Layer*, which represents the core of the architecture and provides a set of domain-independent APIs that are used to access the DKF capabilities; and (iii) *HLA RTI Infrastructure*, which represents the RTI that host the Federation [5] (e.g. PITCH, VT/MÄK, PoRTIco, CERTI). Some application-specific extensions of the DKF can be also introduced (e.g. the SEE-specific ones).

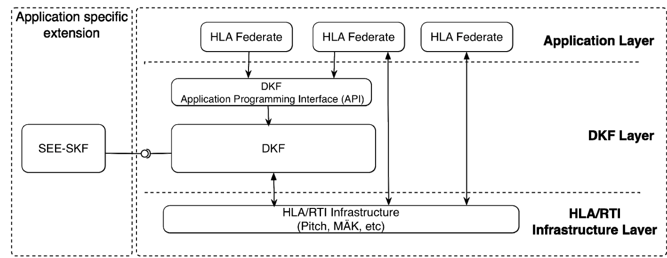

Fig. 1. The architecture of a DKF-based Federation.

The *DKF* is organized into a hierarchy of packages and sub-packages; each of which contains a set of Java classes and interfaces that implement specific functionalities; the main packages are shown in Fig. 2 by using a UML package diagram. In particular, the *DKF* is composed of seven main packages that are independent both of application domains and HLA RTI implementations.

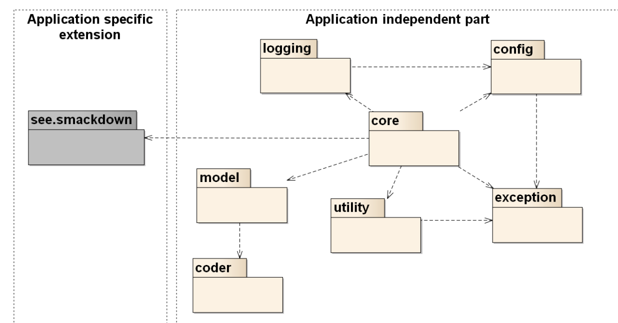

Fig. 2. DKF packages.

The *core* package implements the kernel of the DKF. It includes the fundamental *DKFAbstractFederate* and *DKFAbstractFederateAmbassador* classes (see Fig. 3) that provide the basic functionalities to manage a Federate.

The *config* package contains the collection of classes that manage the configuration parameters provided by a *".json"* file. These parameters include the name of the Federation Execution, the RTI connection details (e.g. IP address, port, etc.), and details about the simulation time.

The *utility* package contains several miscellaneous utility classes, such as time standard conversions (e.g. JulianDate, RJulianDate, etc.) and firewall checking.

The *logging* package defines a set of classes used to track down any problems or errors occurred during the execution of SEE Federates; this information is stored into the *dkf.log* file.

The *exception* package contains some definitions of exceptions that are used for handling dysfunctional events throughout the DKF framework.

The *model* package contains some classes to facilitate publishing, subscribing and the data updating of both *ObjectClass*es and *InteractionClass*es through Java annotations. Two Java annotation classes have been created in

order to manage an *ObjectModel* (*ObjectClass* and *InteractionClass*) instance: (i) *ObjectClassAnnotation*, which defines the annotations that have to be used by the programmer so as to create an *ObjectClass* instance compatible with the DKF; and (ii) *InteractionClassAnnotation*, which specifies the annotations to create and handle *InteractionClass* instances.

The *coder* package contains the classes used for coding and decoding both *ObjectClass* and *InteractionClass* instances.

Finally, the application domain extension *see.smackdown* package, contains some SEE domain-specific classes, which are used by the *core* components of the *DKF* to handle some specific aspects related to a SEE Federation [9] such as transformations among *SEE Coordinate Reference Frames*, the publish and subscribe of *PhysicalEntities*, and the management of the *SISO Space FOMs* [2], [11].

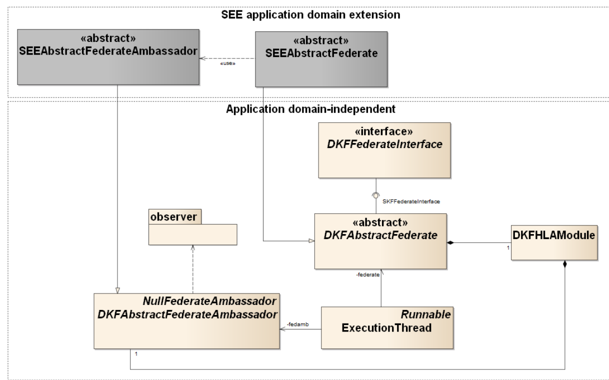

Fig. 3. The architecture of a DKF-based Federate.

2) Federate Behavioral Model

The example architecture of a Federate created by using the capabilities of both the DKF and its SEE-specific extensions is shown in Fig. 3 by using a UML Class Diagram; in the following its main classes are briefly described.

The classes *SEEAbstractFederate* and *SEEAbstractAmbassador*, which are in grey, define the behavior of a SEE Federate, while the classes in yellow belong to the DKF application independent part (see Fig. 2).

The *SEEAbstractFederate* class implements the methods of the *DKFAbstractFederate* class. This latter class provides functionalities to configure and connect/disconnect a Federate to/from a Federation Execution. Moreover, it is worth noting that, in the SEE context, all the Federates are exclusively *time constrained* (can receive Time Stamp Order (TSO) messages) except the *Environment Federate*, provided by NASA and which lead the Federation execution, that is also *time regulating* (can send Time Stamp Order (TSO) messages); the DKF has been thus adapted to handle this situation.

The *SEEAbstractAmbassador* class implements the *DKFAbstractFederateAmbassador* class in order to interact with the RTI services.

Finally, the *ExecutionThread* class handles the execution of a HLA Federate in the simulation environment.

It is worth noting that the *DKFAbstractFederate* class also provides and manages the life cycle of a SEE Federate according to the behavioral model that is shown in Fig. 4 through a UML Statechart diagram. As a consequence, a SEE working team has only to define the specific behavior of its SEE Federate without worrying about low-level implementation details since the DKF manages them. Specifically, the pro-active part of the behavior of a Federate is specified in the "processing and update data" composite state, which is accessed between a TAG (Time Advance Grant) and a TAR (Time Advance Request); whereas, the re-active part of the behavior of a Federate is specified in the "processing interaction" composite state so as to indicate how to handle the RTI callbacks about the interactions/objects that the Federate has subscribed.

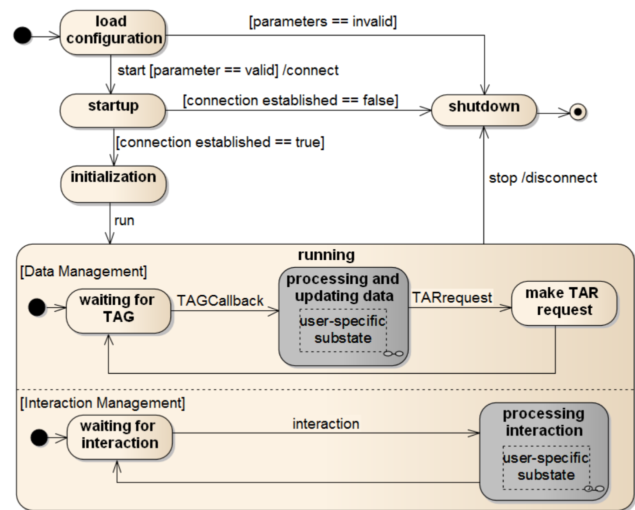

Fig. 4. The life cycle of a SEE Federate.

With reference to the Federate life cycle depicted in Fig. 4, in the *load configuration* state, the DKF loads the configuration parameters from a *.json* file. A transition to the *startup* state happens if the configuration parameters are valid and during the state transition a connection to the SEE Federation execution is performed. Otherwise, if the configuration parameters are invalid a state transition to the *shutdown* state is performed. In this latter state, all the resources engaged by the SEE-SKF classes are de-allocated and the lifecycle terminates. In the *startup* state, the connection status is checked. If the connection is not established the lifecycle ends with a transition to the *shutdown* state.

Otherwise, a transition to the *initialization* state is performed; in this state, the SEE Federate could perform additional operation for exchanging initialization objects before entering the *running* state (and thus the time advancement loop: *waiting for a TAG* → *processing and update data* → *make a TAR*), as an example, the Federate could publish and subscribe some SEE information (e.g. *ReferenceFrames*, *InteractionClasses*, etc.). After that, the time management thread is activated and a transition to the *running* state is performed. The *running* state is composed by two sub-states operating in an AND-decomposition fashion. The *Data Management* sub-state deals with the pro-active part of the Federate behavior through three states: (i) *Waiting for TAG*: the DKF waits for the "TAG (Time Advance Grant) Callback"

from the RTI. When the callback is received a transition to the *processing and update data* state is performed; (ii) *processing and update data*: the "logical time" is updated, the pro-active behavior of the specific SEE Federate defined by the SEE working team is executed, and then a transition to the *make TAR request* state is performed; (iii) *make TAR request*: the DKF requests to the RTI the grant for the next "logical time". The *Interaction Management* sub-state deals with the re-active part of the behavior of the Federate: upon reception of RTI callback related to subscribed elements, a transition to the *processing* state is performed where the received information is handled.

When the simulation ends a transition from the *running* state to the *shutdown* state is performed and, during the state transition, the HLA Federate is disconnected from the RTI.

IV. Exploiting The SEE HLA Starter Kit

This section presents a case study concerning the development of a HLA Federate by using the SEE HLA Starter Kit in the context of the SEE 2015 project. The architecture and behavior of the developed and experimented SEE HLA Federate are described along with the feedback coming from the experimentation.

A. *The Development Process based on SEE-SKF*

The process to build a Federate from scratch by using the SEE-SKF is composed by the following four main steps:

(i). Build a *model* of the Federate that specifies: the *objects* that the Federate manages (as specified in the FOM), the *attributes* of these objects and the *coders* to handle such attributes. It is possible to use the set of basic coder provided by the SEE-SKF or simply implement new coders by using the SEE-SKF classes; other available coders can be also exploited;

(ii). Build a concrete Federate that specifies the *behavior* of the *model* defined at step (i). It is required to extend the *SEEAbstractFederate* abstract class provided by the SEE-SKF and implement three methods according to the Federate life-cycle that is provided and completely managed by the SEE-SKF (see Fig. 4), specifically: (a) a method for initialization operations before entering in the "running state" (a *configureAndStart()* method); (b) a method for specifying the pro-active part of the behavior of the Federate (*doAction()* method) and that is executed between a TAG and a TAR; (c) a method (*update* method) that specifies the re-active part of the behavior of the Federate, i.e. how to handle the RTI callbacks about the interactions/objects that the Federate has subscribed;

(iii). Implement the Federate Ambassador. This step requires extending the *SEEAbstractFederateAmbassador*; typically, since no specific implementation is required, the child class has only to define its constructor which in turn calls the parent one: all the typical Ambassador's features are provided and managed by the SEE-SKF;

(iv). Implement a *main* class so as to instantiate and run the developed Federate.

In the following, after presenting the reference simulation scenario, the above sketched process will be exemplified with respect to the development of a Federate in the context of the SEE Project [9].

B. *The reference Simulation Scenario*

The reference simulation scenario of the SEE Project (see [2], [11]) concerns a human settlement called "Moonbase" composed of scientific equipment, storage buildings, rovers and other elements to allow astronauts to live and work on the Moon.

The Modeling & Simulation Group (MSG) at Brunel University London has participated in the SEE Project since 2013. The group has investigated issues concerning the development and standardization of distributed simulation for industry and healthcare [12] as well as hybrid federations consisting of real-time, discrete-event and agent-based simulations [11].

The main issue that arose from the SEE 2014 event was the complexity of the development. The students based their work on previous code developed by the group. However, the broad knowledge base of domain specific knowledge, distributed simulation (both Federate development and RTI interfacing) and the SEE event scenario still presented a major challenge due to the range of possible implementation approaches and the lack of clear development guidelines and tutorials.

In the SEE 2015 entry participating was restricted to one undergraduate student. His task was to develop a new agent-based mining simulation that simulates one or more small excavators working across the lunar surface. The simulation is designed to work with an astronaut Federate and a UAV Unmanned Aerial Vehicle) Federate.

C. *The Escavator Agent-based Simulation*

REPAST SIMPHONY [8] is a free and open-source agent-based simulation environment [3]. A *REPAST* agent-based simulation is created by using the *ContextBuilder* interface. In this class, the environment, the initial number of agents (and types/classes) that are located in the environment, etc. are specified. The attributes and methods of each agent are specified in an agent's class. Each agent interacts with other agents and the environment via their methods. Time is managed in a *REPAST* simulation by the scheduler. A method can be annotated as being scheduled. This will include the frequency and priority that the method occurs. When a *REPAST* simulation runs, the simulation environment enters a cycle that calls the scheduler. The scheduler then runs the methods in priority order according to their frequency, so advancing the simulation until it reaches some terminating condition.

The ultimate goal of the excavator agent-based simulation was to explore how excavator "robots" could self-organize in the coordination of the extraction of lunar regolith materials and the degree to which *REPAST* could facilitate the study of these algorithms. As this paper focuses on how the SEE-SKF

was used to simplify the implementation of a Federate, a simple version of the agent-based simulation is presented. In this example there is a single excavator that explores its environment in a simple "scanning" pattern. It uses a map populated by a UAV with a magnetometer. The UAV slows "flies" overhead detecting potentially interesting minerals (modelled by a reading of 0 for nothing, 1 for something). The UAV periodically broadcasts the results of its on-going survey to the excavator and the excavator updates its local map (ten sets of readings each update). When the excavator reaches a mineral, it mines it and adds it to its hopper that carries the excavated regolith. Once the hopper is full the excavator returns to its origin point and deposits the regolith material in a pile. The now empty excavator returns to where it left off and continues mining.

The agent-based simulation consists of three main classes: the *JExcavatorsBuilder*, *Excavator* and *Mineral* class. *JExcavatorsBuilder* implements the *REPAST ContextBuilder* interface to create the simulation environment. It does this by first creating a continuous space and then superimposing a grid for the excavators to move around. The grid is then populated by an Excavator at 0,0. An Excavator has several variables that specify where it is currently located on the grid (*pt*), the amount of regolith carried (*cargo*), and if it has decided to return to the origin point to offload its regolith (*returnOrigin*).

When the simulation begins it calls all *step()* methods in its agents. In this example, the single excavator *step()* method is called. This first simulates the interaction with the UAV by generating the next ten magnetometer readings and updating the map by calling *updateMap()*. This populates the grid with zero to ten new Minerals at x,y points (i.e. UavX1, UavY1, MagReading1, etc., where *MagReading* is either 0 or 1 to indicate the presence of a mineral deposit). The excavator then reads its own location on the grid and checks to see if its cargo limit has been reached. If it has, it sets a Boolean *returnOrigin* to TRUE. If *returnOrigin* is FALSE, the *moveLinearly(pt)* method moves the excavator along to the next point, checks the map to determine if there is anything to mine and, if there is, mines it (adds 1 to the cargo and deletes the mineral from the map). If *returnOrigin* is TRUE, then this method excavator first moves the excavator one point at a time along its X axis to 0 and then its Y axis to 0 to reach the origin. The excavator then "dumps" its cargo and retraces its steps back to where it left off, again one point at a time.

```
1. //@ScheduledMethod(start=1, interval = 1)
2.      public void step()  {
3.              updateMap();
4.              checkCargoLimit(cargo);
5.              moveLinearly(grid.getLocation(this));
6.   }
```

D. *Using the SKF to Develop the Excavator Federate*

The above description of the simple excavator focuses on a single excavator agent. The mining operation may be also of interest to other simulations (e.g. an astronaut who takes away mined materials for processing). To create a Federate based on the above introduced agent-based simulation the SEE-SKF main steps have been followed. In step (1) a FOM that describes the input and output of the simulation was defined. In this case the FOM represents the single Excavator object with *ExcavatorX* and *ExcavatorY* representing the current coordinates of the excavator, and *PileNumber*, representing the number of minerals in the regolith pile. All are *HLAinteger32BE* datatype. To begin the creation of the Federate, the Excavator class has been annotated to match the FOM as follows:

```
1. @ObjectClass(name = "PhysicalEntity.Excavator")
2. public class Excavator{…
```

To create the I/O from the simulation to the rest of the Federation, the Excavator class was augmented with attributes and coders. For example, to enable the sharing of the X, Y coordinates of the excavator the following attributes and coders have been added to the declarations:

```
1. @Attribute(name = "ExcavatorX", coder = HLAinteger32BECoder.class)
2. private Integer ExcavatorX;
3. @Attribute(name = "ExcavatorY", coder = HLAinteger32BECoder.class)
4. private Integer ExcavatorY;
```

At the end of the *step()* method described above, the two calls

```
1. setExcavatorX(getPointX());
2. setExcavatorY(getPointY());
```

have been added to update the current position of the excavator. Similar attributes and coders for the other attributes described in the FOM have been added.

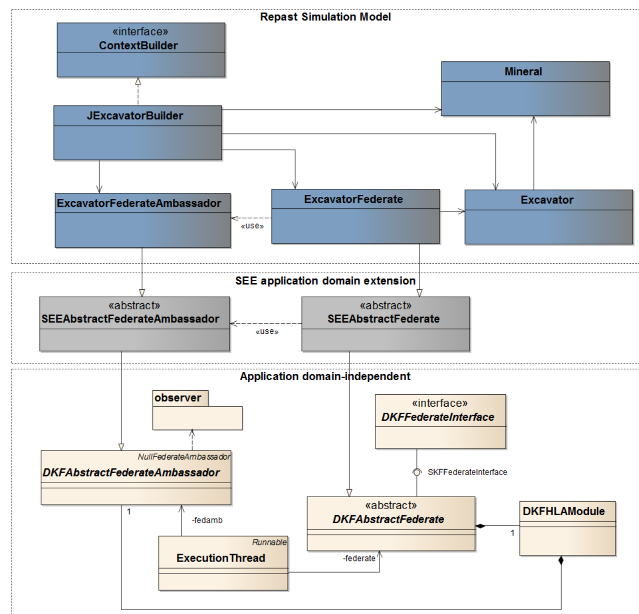

Fig. 5. The architecture of the Excavator Federate.

In step (2), the *SEEAbstractFederate* class has been extended to create the *ExcavatorFederate* class. Within the *ExcavatorFederate* class the *configureAndStart()* method remained unchanged (i.e. it reaches the JSON config file and starts the Federation). The *doAction()* method is shown below.

```
1. protected void doAction() {
2.      for (Object obj : RunState.getInstance().getMasterContext()) {
3.              if(obj instanceof Excavator) // update the excavator on RTI
4.                      ((Excavator) obj).step();
5.      super.updateElement(obj); } }
```

This method advances the agent-based simulation by first obtaining the current state (context) of the simulation, finding all agents (objects) and then "manually" running the *step()* method in the agents. In this example, the single excavator agent's *step()* method is executed. It then calls *updateElement(obj)* to output the new state of the excavator Federate's attributes.

Step (3) simply extended the *SEEAbstractFederateAmbassador* class with the *ExcavatorFederateAmbassador*. Step (4) was unnecessary, as the simulation had already been developed. The only addition to these steps was that of the *ExcavatorFederate* and *ExcavatorFederateAmbassador* to the context (*JExcavatorsBuilder*) to include them in the scope of the agent-based simulation. The overall class diagram is shown in Fig. 5.

V. DKF Qualitative Assessmnet

The challenge presented to the Brunel undergraduate team was how to create a simulation of a set of self-organizing excavators and their mining operation. The first task was to build the simulation in *REPAST*. The team had never done any kind of simulation before but had some experience in Java programming. However, *REPAST* has reference examples and documentation that could be used to support the teams' development. The second task was to then implement a Federate based on the agent-based simulation. In 2014's SEE event this proved to be an extremely challenging task for the team of that year, despite the support of the MSG team. The main problem was the lack of support material that the team could use. Very little existed at the time apart from Fujimoto's text book [3], Moller's introduction to the HLA [7], and "hints" in key articles on HLA issues (e.g. time management). The experiences of using the *DKF* and its associated process had a great impact on the development time of the Federate as much of the HLA complexity was hidden away. The first attempt to understand the *DKF* and *SEE-SKF* was to produce a simple astronaut Federate. This essentially allowed a user to move a point across a space using a keyboard. The simulation just produced the coordinates of the astronaut. Following the SEE-SKF process, the team analyzed the data produced and required by the simulation (e.g. coordinates) to create the FOM and then used the SEE-SKF attributes and coders to map the FOM to the input/output in the simulation. The *SEEAbstractFederate* class was then implemented with its methods: the *configureAndStart()* method to run, the *doAction()* method to move the astronaut, and the update method to send the current coordinates of the astronaut. Finally, the Federate Ambassador was implemented and the main class was made available.

Developing the *REPAST* Federate proved more challenging as it was not clear at first how the above mapped to the environment and agent classes of the simulation. The previous section presented how this was achieved. The only unsatisfactory aspect of the implementation was the delegation of *REPAST* time management to the Federate Ambassador. *REPAST* has excellent time management facilities and this would have been better if the Federate Ambassador had been coordinated with *REPAST* time management. We expect this issue to be resolved in future work. To test the simplicity of using the SEE-SKF, the Excavator was linked to the UAV simulation developed by Liverpool University that hovers over the excavator grid. The UAV is equipped with a magnetometer that takes readings from the surface to identify interesting areas for excavation. To add this to the Excavator we assume that the UAV will produce ten sets of readings each update. *UavX1*, *UavY1*, *MagRead1*, etc. were then added to the FOM under a UAV class to allow the Excavator to subscribe to the attributes. Attributes and coders were added to the Excavator agent along with a *GetMap()* method that takes the current UAV updates and adds these to the excavator's map. Once the FOM had been agreed with Liverpool it took the team a remarkably short time update the agent and get the Excavator/UAV distributed simulation up and running.

VI. DKF Quantitative Assessment

This section presents a quantitative analysis about the quality of the code produced by using the DKF and aims at highlighting the benefits provided by its exploitation in the SEE project.

Software complexity is a primary topic in Software Engineering and has involved many researchers over the years. Many metrics related to various constructs like class, coupling, cohesion and inheritance have been proposed, and they play a fundamental role to measure the quality of source code. To evaluate the complexity of the source code of an SEE-SKF based HLA Federate, six standard metrics, which are proposed by various researchers, have been considered [16]: (i) *SLOC (Source Line of Code)*, which represents the number of executable statements; (ii) *CCM (Cyclomatic Complexity Metric)*, which is a metric based on graph theory that represents the number of linearly independent paths through a program's code; (iii) *HCM (Halstead Complexity Metric)*, which measures the logic volume of the code. It is calculated on the count of the operators and operands. The operators are symbols used in expressions to specify the manipulations to be performed, whereas the operands are the basic logic unit to be operated; (iv) *NP (number of package)*, which is the number of packages; (v) *NF (number of function)*, which represents the total number of functions; and, (vi) *NC (number of classes)*, which represents the number of classes. Typically, the lower are the values of these metrics the lower is the complexity of the source code and thus higher should be the code compactness, readability and reliability. These six metrics are evaluated comparing the source codes of the same HLA Federate called *UNICOM* [2] on the same functionalities (see TABLE I.). More in detail, one source code is based on the SEE-SKF whereas the other one is that produced by the *Pitch Developer Studio* [7]. The metrics have been calculated by using the *Google CodePro AnalytiX* tool, which is an application, developed by Google Inc., that allows developers to perform code measurement and comparison with user-defined programming standards and that is used by several large organizations, ranging from aerospace/defense to automotive/transport companies, to control their programming process.

Although the *DKF* framework, and its domain-dependent extension SEE-SKF*,* does not cover all the IEEE 1516-2010 functionalities that are instead covered by the *PITCH Developer Studio*, the results reported in TABLE I. show that

the source code of an HLA Federate created by using the *DKF*/*SEE-SKF* is easier to manage and maintain compared to the same code produced by the *PITCH Developer Studio*. Moreover, some classes of the second one tool have *CCN* value more than twenty-two; as a consequence these classes are difficult to manage/extend by programmers (see [16] for a discussion).

TABLE I. METRICS AT PACKAGE LEVEL

Metric		UNICOM Federate	
		SEE-SKF	Pitch Developer Studio
NP		9	8
NC		17	72
NF		94	784
SLOC		744	6186
CCM (average)		1,20	1,63
HCM	Number of distinct operators (n_1)	22	38
	Number of distinct operands (n_2)	312	1454
	Total number of operators (N_1)	1337	4150
	Total number of operands (N_2)	513	13217
	Software length (N)	1850	17367
	Software vocabulary (n)	334	1492
	Volume (V)	$1{,}584*10^4$	$1{,}831*10^5$
	Level (L)	0,1495	0,4784
	Difficulty (D)	47,13	172,71
	Programming Effort (E)	$7{,}469*10^5$	$3{,}162*10^7$
	Error Estimate (B)	5,28	61,03
	Programming Time (T)	$4{,}149*10^4$	$1{,}756*10^6$

VII. CONCLUSION

HLA is undoubtedly one of the most mature and popular standard for distributed simulation. Due to its capabilities to enable the interoperability and reusability of distributed simulation components, it is increasingly exploited in a great variety of applications in both military and civil domains. However, the development of full-fledged simulation models, based on HLA, is still a challenging task. In this context, the paper has presented the *HLA Development Kit*, a general-purpose, domain-independent toolkit that aims at easing the development of HLA-based simulations by providing a software framework (the *DKF*), with related documentation, user guide and reference examples. The effectiveness of the *DKF* has been exemplified in the context of the Simulation Exploration Experience (SEE), an international project lead by NASA and which involves several U.S. and European Institutions in the distributed simulation of a “Moonbase”.In terms of developing educational resources for HLA development, the *DKF* presents a solid foundation for future expansion. The SEE event is exciting in that students can create a wide variety of simulations and take part in an international project. The *SEE-SKF* is therefore an example of how the *DKF* can be extended to be domain specific. Future work on *DKF* includes the further development, testing and evaluation of the *DKF* and its domain specific extensions so as to also provide interesting education and research resources to easily develop distributed simulations in various application domains.

ACKNOWLEDGMENT

The authors would like to thank Edwin Z. Crues (NASA JCS) for his precious advice and suggestions in the development of the *HLA Development Kit*. A special note of thanks goes also to all the NASA staff involved in the SEE Project: Priscilla Elfrey, Stephen Paglialonga, Michael Conroy, Dan Dexter, Daniel Oneil, to Björn Möller (PITCH Technologies), and to all the members of SEE teams.

REFERENCES

[1] J. Banks, J.S. Carson, B.L. Nelson, and D.M. Nicol, Discrete-Event System Simulation, 5th Ed., Prentice Hall, 2009.

[2] A. Falcone, A. Garro, F. Longo, and F. Spadafora, Simulation Exploration Experience: A Communication System and a 3D Real Time Visualization for a Moon base simulated scenario. In *Proc. of the 18th IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (ACM/IEEE DS-RT)*, pp. 113-120, IEEE Computer Society, 2014.

[3] G. Fortino, A. Garro, W. Russo, From Modeling to Simulation of Multi-Agent Systems: an integrated approach and a case study. In *proc. Of the 2nd Multiagent System Technologies*, pp. 213-227, 2004.

[4] R. M. Fujimoto, Parallel and distributed simulation systems, John Wiley & Sons, 2010.

[5] IEEE Standard for Modeling and Simulation (M&S) High Level Architecture (HLA): 1516-2010 (Framework and Rules); 1516.1-2010 (Federate Interface Specification); 1516.2-2010 (Object Model Template (OMT) Specification).

[6] MÄK VR-Forces, [online], available http://www.mak.com/.

[7] B. Möller, The HLA tutorial v1.0, Pitch Technologies, Sweden.

[8] M. J. North, N.T. Collier, J. Ozik, E. Tatara, M. Altaweel, C.M. Macal, M. Bragen, and P. Sydelko, Complex Adaptive Systems Modeling with Repast Simphony, Springer, 2013.

[9] Simulation Exploration Experience (SEE) project, [online], available http://www.exploresim.com/.

[10] S.J.E. Taylor, P. Fishwick, R. Fujimoto, E. Page, A. Urhmacher, G. Wainer, Panel on Modeling & Simulation Grand Challenges. In *Proc. of the 2012 Winter Simulation Conference (WSC)*, pp. 1-15, Association for Computing Machinery Press, New York, 2012.

[11] S.J.E. Taylor, N. Revagar, J. Chambers, M. Yero, A. Anagnostou, A. Nouman, and N.R. Chaudhry, Simulation Exploration Experience: A Distributed Hybrid Simulation of a Lunar Mining Operation. In *Proceedings of the 18th IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (ACM/IEEE DS-RT)*, pp. 107-112, 2014.

[12] S.J.E. Taylor, S.J. Turner, N. Mustafee, and S. Strassburger, Bridging the gap: a standards-based approach to OR/MS distributed simulation, *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, Vol. 22(4), Article 18, 2012.

[13] The Forwardsim HLA Toolbox for MATLAB, [online], available http://www.forwardsim.com/products/hla-toolbox/.

[14] O. Villimann, CTO Project, Documentation, HLA Framework. Danish Maritime Institute, 1999.

[15] Y. Xie, Y.M. Teo, W. Cai, S.J. Turner, Towards grid-wide modeling and simulation, 2005.

[16] S. Yu, S. Zhou, A survey on metric of software complexity. In *Proc. of the 2nd IEEE International Conference on Information Management and Engineering (ICIME)*, pp. 352-356, IEEE Computer Society, 2010.