

GaliLEO: a simulation tool for LEO satellite constellations ^{*†}

Laurent Franck
ENST/TéSA
Toulouse (FR)
Laurent.Franck@enst.fr

Francesco Potortì
CNUCE-CNR
Pisa (IT)
F.Potorti@cnuce.cnr.it

Abstract

We present GaliLEO, a simulator for the transmission of connection-oriented traffic over a constellation of LEO/MEO (Low / Medium Earth Orbit) satellites. Its scope is limited to the satellites and the stations accessing them, without any modelling of the terrestrial network, but inside this scope the goal is to study the performance of satellite-based communication networks from as many as possible points of view at the network level. Typical applications include simulation of access techniques, routing policies, fault management. The simulator is written in Java, and it makes use of dynamic loading to easily integrate user-written modules. A draft manual is available, and a preliminary version of the program has been published at the end of 2000.

1 Introduction: the newborn and the world around it

The motivation behind GaliLEO's conception emerged during an exchange of ideas among some members of the European COST 253 Action¹, a forum where researchers from all around Europe periodically meet to address issues related to LEO constellations of communication satellites. The point was made that none of the commercially or freely available simulation tools was reasonably usable as a generic simulation tool for LEOs.

^{*}The final, definitive version of this paper has been published in SIMULATION, Volume 78, Issue 9 of September, 2002 by Sage Publications Ltd, All rights reserved. ©The Society for Modeling and Simulation International (Simulation Councils Inc), 2002. It is available at: <http://online.sagepub.com/>

[†]This work was partially supported by the CNR (Italian National Research Council) under the 5% Multimedia Programme.

¹COST stands for *Cooperation in the field of Scientific and Technical research*, see <URL:<http://www.eeng.brad.ac.uk/Research/cost253/>> for information on COST 253, which ended in June 2000.

This seminal discussion later led to the initial design of the GaliLEO architecture, which was the outcome of a collaboration between two institutes where researchers had already had experiences in developing special purpose simulators. The objective of GaliLEO is twofold. First, it has to provide a simulation framework strongly focused on satellite constellations in order to relieve the user from all the burden of tedious programming about satellite/constellations characteristics (orbital mechanics, handover management, access methods). Second, the framework has to be open in order to allow modifying the implemented mechanisms as effortlessly as possible. These two objectives call for an object oriented approach where the simulator is made from building blocks. Some of these blocks may be replaced (although their location in the framework is fixed) at the user's convenience, achieving customisation of the simulation model. Other blocks may not be changed, taking part in immutable behaviour of the simulator. Achieving adaptivity while keeping a tight focus on satellite constellations was the main challenge during the design of GaliLEO.

1.1 Other network simulators

The literature is rich in simulation papers. Many of them do not mention the tools used to obtain the results. Most of the simulators that are explicitly mentioned belong to the discrete-event family, i.e. they model a certain system as it evolves over time using a representation where the state variables only change at a countable number of points in time. Events occur at these points: they are defined as instantaneous occurrence which may change the state of the system.

In the following, we very briefly describe four simulators that are often cited in the literature, and we mention their perceived drawbacks with regard to large-scale, network-level simulations of satellite constellations. Two problems common to all the cited simulators, apart from *ns*, are that they are generally expensive for small research groups and

for students, and that their development model is not open.

OPNET (OPTimised Network Engineering Tool) is a simulation tool, based on state machines, for analysing communication networks by using models[9]. Opnet models are specified in terms of objects, each with a configurable set of attributes that can be specified either by a graphical process, an associated text file, or as variable parameters in the simulation description file. The specification of the models is organised into a hierarchy of four different levels, called network, node, process, and parameter. Once the model is specified, Opnet generates a simulation program written in the C language. Opnet is a very general simulator that can be used to simulate satellite constellations. Its general purpose character is however a weakness when studying complex configurations, because Opnet needs to identify anything as an event, which makes ad-hoc optimisations impossible or very difficult. In [4], a test case is made showing an impressive speed difference in excess of 100:1 between the specialised simulator Fracas and Opnet. Generally speaking, we expect a big performance difference between any specialised simulation tool and general purpose ones like Opnet and Bones.

BONeS (BlockOriented Network Simulator) is a simulation system for studying communication network models[14]. *BONeS SatLab* is a software package for the design, animated visualisation and analysis of satellite-based communication systems. It provides multiple animated views of global satellite systems, uplink, crosslink and downlink analysis, jamming, and adjacent satellite interference analysis representation of earth stations. Three types of simulation can be performed: positioning, design and communication simulation. The positioning simulation lets the user model analyse and animate various configurations of satellites, fixed earth stations and moving stations. The design simulation automatically performs simulation for a range of parameter values. For example, the user can determine the interference of communication satellite systems on other communication systems or the probability of interference between satellite systems for selected frequencies. The communication simulation is used to track the source and route of data packets and to determine their best route based on relative distance, velocity, angle, visibility, traffic congestion, and interference between two or more nodes. While very flexible, Bones is focused on the radio and telecommunications aspects, or else packet-level simulation; therefore it is not particularly efficient for simulating networking aspects by running long simulations with a great number of connections.

ns-2 (Network Simulator) is a free TCP simulator at the packet level. It is not usable for our purposes, because it provides no support for dynamic network topologies or MAC level simulation, apart from a CSMA/CD module targeted at LANs. It is available at

<URL:<http://www.isi.edu/nsnam/ns/>>.

RESQ (RESearch Queuing package) is a software package developed at IBM Research for defining and solving extended queueing network models[12, 13]. Resq provides a numerical solution component, *QNET4*, which uses the convolutional algorithm for product form networks, and a simulation component, named *APLOMB*. Resq is especially strong in the statistical analysis of the simulation output and the computation of appropriate simulation run lengths. Its main limit is that it is specially focused on systems that can be naturally be modelled with a queueing system.

1.2 GaliLEO's ancestors

The basic idea and most of the concepts regarding the connection setup and the channel access from a ground station were developed at CNUCE, an institute of CNR in Pisa (IT), as a consequence of the inadequacy of the locally developed Fracas [4] simulator for the study of LEO networks. The routing concepts and the details of the architecture that form the glue of the actual implementation of GaliLEO come from the experience of the LeoSim [7] simulator, developed originally at the Brussels University (BE), and currently at ENST in Toulouse (FR), and the initial specifications of SimToc [6], designed at CNUCE. Interestingly, while having very different scopes and objectives, all of these simulators took an object oriented approach to implementation, principally as a means to ease extensibility. GaliLEO aims to be a general purpose, customisable tool, freely available for the whole satellite community.

Fracas (FRAmEd Channel Access Simulator) is essentially a command line driven emulator whose time advances in fixed length steps, usually of the same length of a frame of the protocol under study. While very fast and very well suited to the study of access protocols for GEO systems, it cannot be adapted to LEO systems. Its heritage consists of the concepts behind statistics collection and manipulation.

LeoSim, the most important of GaliLEO's ancestors, is an event-driven, continuous time simulator accessed through a graphical interface. It has been developed at ENST (FR) in order to study link state routing algorithms for LEO satellite constellations. LeoSim provides statistics on the number of call requests, the call block probability, and the cost introduced by maintaining the link state database; shortest path routing, handover management and elaborate routing signalling are implemented. Its design approach and its core simulation engine have been transported into GaliLEO. The object model has been streamlined and enhanced in order to add support for access methods, beam management and enhanced call signalling. Several hooks have also been placed in the network model in order to ease the future evolution towards connectionless traffic. Its most prominent limitation is that access to the satellite is not sim-

ulated at all. This makes it impossible to simulate methods for accessing the satellite channel from the earth stations.

From *SimToc*, the other ancestor from CNUCE, GaliLEO took the global architecture, the idea of the up-down link between the ground station and a satellite of the constellation, and the way a connection is set up and modified. *SimToc* has never gone past the design stage.

Consim [1] is a simulator developed at CSELT (IT) for evaluating the performances of constellations of communication satellites affected by different types of failures. *Consim* will be integrated by results into GaliLEO. By “integration by results” we mean that the two simulators are kept separate, and the results of *Consim* are used by GaliLEO. This is a simple way to program interaction between two simulators that were written separately, while minimising the coupling needed between the different teams responsible for the programs. However, this approach is only feasible when the two studies cover aspects which are not interdependent. In this case, *Consim* runs a failure model to produce a list of failure events occurring during the constellation lifetime, each one tagged with the type of failure and the time of occurrence. Since the fault occurrence is independent of the traffic generation, *Consim* needs no feedback from GaliLEO, and the data exchange between the simulators can be unidirectional. In practice, the list of failure events is provided during the simulation initialisation phase and then used to feed GaliLEO’s simulator engine in order to trigger the right fault managers at the appropriate time.

2 Architecture

Figure 1 highlights the three-layer architecture of GaliLEO.

The first layer (*Simulation engine*) defines both the structure of *components*, which are the blocks from which GaliLEO is built, and the way they behave and interact. The components’ dynamics is the job of the scheduler that runs the components and defines a message passing structure for inter-component communication. Thus the first layer constitutes a generic infrastructure for a discrete-event, message-passing simulator.

The second layer (*Core modules*) implements the *network model* of GaliLEO, by defining the scope of GaliLEO possibilities through the definition of a set of classes. These classes (called *templates*²) also specify the rules for using the network model and creating custom components. This layer is not customisable per se, and in fact is the core of GaliLEO’s functionalities. Experimenters needing to implement their own set of modules should be well acquainted with the network model.

The third layer (*Custom modules*) is the set of modules which are dynamically loaded from a library, including both

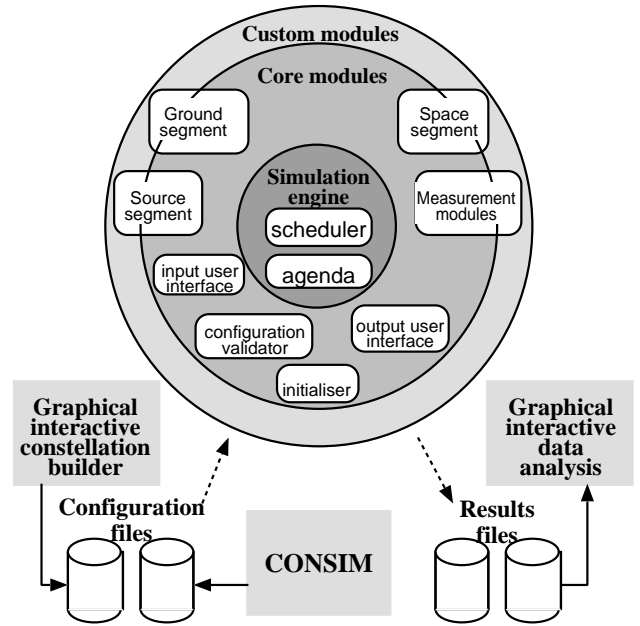


Figure 1. GaliLEO architecture.

standard and *custom components*. Standard components are components shipped with GaliLEO. Custom components are developed (possibly based on standard components) by the user of GaliLEO in order to tailor the simulator to its needs. This layer is where ad-hoc built modules are integrated in GaliLEO. Examples include modules defining the behaviour of actual routing algorithms, channel allocation methods, traffic generators, call admission control policies, and so forth.

To summarise the relationship between the second and third layer, layer two defines what are the general characteristics of, say, a channel allocation method (in terms of what are the provided services) while layer three defines actual channel allocation methods.

2.1 Components as building blocks

GaliLEO is extremely modular, because it aims at providing a simulation framework where one plugs in a locally developed, e.g., routing algorithm, and evaluates the resulting behaviour. The basic module is called a *component*, which is a class of Java [3] objects that provide methods to duplicate, initialise and start themselves after creation. *Initialisation* may be based on the presence of other components in the system. *Starting* a component is done after initialisation. This usually makes sense only for *entities* (modules with a special processing capability), which are described below. GaliLEO comes with a small collection of standard components, which are meant to be used as-is or replaced with custom ones. Hopefully, GaliLEO’s library

²They are unrelated to C++ templates.

of standard components will grow with time.

Standard and custom components are built upon templates, which are Java abstract classes used to provide an API for the development of components. Providing an API has some shortcomings with respect to providing an extension language; it is generally more difficult to program in Java than in a dedicated extension language, which can be limited to provide only special constructs and can be well insulated from the details of the simulator core. However, using an API is a far easier and more flexible approach, and certainly more efficient in terms of resource usage.

Components contain both code and data. After initialisation, a component can either live as a passive element whose methods are called by the system and other components, or behave as an independent piece of code. This latter special kind of component is called an *entity*. Entities run concurrently with the rest of the system and other entities, by using the communication and scheduling facilities provided by the simulation engine.

The GaliLEO network model described later is therefore a collection of inter-operating components.

3 The simulation engine

The simulation engine comes from LeoSim, and includes the scheduler and the agenda.

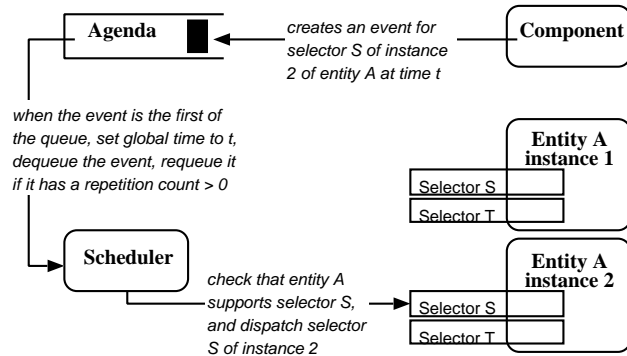


Figure 2. Creating, scheduling, and consuming events.

Any module inside GaliLEO can generate an event by calling the `scheduleAction` method of the *scheduler*, which creates a pending event. The arguments of this method are the delay after which the event should be triggered, a repeat count, and the action triggered by the event. The scheduler organises the pending events in a structure called an *agenda*, which is conceptually a queue where the events are kept sorted according to the time when they should be triggered. The exact implementation of the

agenda is customisable, to allow experimentation, easy upgrading, and platform-specific optimisations. Currently, a simple-minded delta list is implemented, together with a more sophisticated calendar queue implementation.

A *delta list* is a structure allowing basically two operations, namely insertion of a random element and extraction of the smallest element. The implementation consists of a linked list where an event is inserted in order, so the extraction consists simply of extracting the first element of the list. Inserting an element needs scanning the list from the beginning, so the insertion time is $O(N)$, N being the number of pending events, which is a measure of an event driven simulation's size. In order to perform big simulations like those we are planning, all computation related parameters, prominently memory consumption and execution speed, should scale not much faster than the simulation's size. Unfortunately, if the insertion of each event in the ordered list takes a time proportional to the total number of events, the total overhead introduced by event scheduling in a simulation will be $O(N^2)$.

A *calendar queue* is a structure allowing the same operations as those provided by a delta list, but using a more complex data structure, consisting of an array of linked lists. It is possible to make an analogy with a calendar, where for each day one writes down zero or more appointments, ordered by the time of the day. Finding the day where the appointment should be written is $O(1)$, and inserting it in the queue of that day is $O(n)$, n being the number of events (appointments for that particular day). Two parameters must be set for a calendar queue, that is the slot size, which is the length of the day in the calendar analogy, and the number of slots. It has been shown in [5] that the optimal number of slots is $O(N)$. By changing dynamically the number and the size of the slots depending on N we obtain a *dynamic* calendar queue, for which empirical evidence has been given in [2] that insertion and extraction times are $O(1)$. Currently, GaliLEO implements a *static* calendar queue, whose parameters are read from the configuration.

The action triggered by an event is defined by a *selector* and a list of arguments to it. A selector is an entry point in a module, that is, a method which possibly accepts arguments. When an event is triggered, the associated selector is called, and the relative list of arguments is passed to it. This simple message passing mechanism allows asynchronous communication between components. More precisely, any piece of code inside GaliLEO can generate an event, and thus send a message, but only entities can have selectors, and thus be awakened by the scheduler and receive a message.

The scheduler is the heart of the simulator. After the initialisation phase, the simulation consists of a loop running inside the scheduler, which just removes the first event from the queue, advances the system's time to that of the event,

and calls the selector specified therein, with the appropriate argument list. When the selector is finished, it returns to the scheduler. The loop ends when there are no more events in the queue, a special stopping event is encountered, or when manually stopped by the operator.

Since it is anticipated that GaliLEO will go distributed in the future, the scheduler is customisable, to allow experimentation and local customisations of distributed scheduling criteria. Currently, a simple serial scheduler is available, which is the normally used one.

Moreover, if the first selector creates any events whose time is less than the time of the second event, these events (recursively) should not trigger any change of state on which the second selector relies. In order to ease this requirement, a second argument is passed to the `isSafeWith` method, which is the time of the second event.

This mechanism is far from being automatic, but in practice it can allow some parallelism for carefully crafted components that have been written by the same programmer. For example, computing a route is a time consuming task (it involves usually a shortest path algorithm). If two successive routing events are present in the agenda, they could be launched concurrently provided that they do not have to be performed in the same satellite.

4 The network model

The second layer of the architecture depicted in Figure 1 defines the basic capabilities of the simulator as far as the modelling of the communication network is concerned. The relevant modules are the Source, Ground and Space segments. Each is a collection of components and templates. Custom and standard components are instantiation of templates, and occupy the third layer of the architecture. GaliLEO will initially ship with a small set of standard components, and a manual describing the API for building custom ones.

4.1 Assumptions and definitions

Many components in GaliLEO are meant to describe real objects in the satellite network. We describe the main concepts used when describing the network, and when there is a direct correspondence between a concept and a component, we will indicate the name of the component in mono-spaced face between brackets, like in `[Satellite]`.

We define a *cell* as the area of the earth illuminated by a satellite spot beam. A *footprint* is the whole coverage area of a satellite `[Satellite]`, i.e. it is the sum of the areas covered by its spot beams. An *overlap area* is the area in which a *ground station* `[Station]` (i.e. a single subscriber or a concentrator) can receive a signal with an acceptable

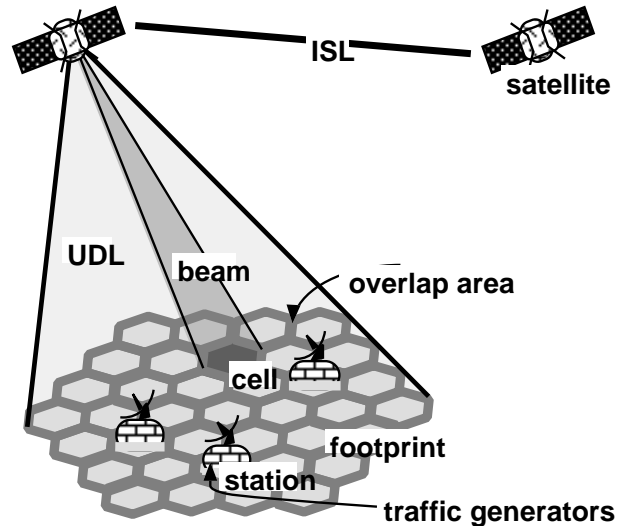


Figure 3. Names of some objects used in the simulator.

power level from more than one adjacent spot beams. A *UDL* (Up-Down link) `[Udl]` is the aggregation of all spot beams pertaining to the same footprint; it has a fixed capacity, and is unidirectional. A *beam* `[Beam]` is the communication medium between a satellite and a spot on the ground. A beam has a variable capacity which cannot exceed the capacity of the UDL the beam belongs to. A *node* `[Node]` of the network is any station or any satellite. Satellites have multi-beam antennas for up-link reception and down-link transmission, and are connected to each neighbouring satellite by means of an *ISL* (inter-satellite link) `[Isl]` which is a unidirectional link.

A *connection* `[Connection]` is a virtual communication path between a *source* and a *destination*, which are normally different stations. A connection can be created, modified by changing its characteristics `[Resources]`, and torn down. It is assumed to be full-duplex, composed by a forward and a return channel `[UniConnection]`, where the forward channel is intended to be from source to destination, and the return channel from destination to source.

The procedure supporting the transition from one connection state to another is implemented by a call signalling protocol `[CallSignaling]`.

4.2 Call signalling

A *call generator* `[CallGenerator]` defines when a connection starts, between which endpoints, and how and when it is modified and torn down. It can be associated with a *packet generator*, which produces the packet traffic running over a connection, for simulating connection-oriented

traffic. It is envisaged that, in a future, GaliLEO may be able to support also traffic generators that create connectionless traffic. In the following we will only consider connections and connection-oriented traffic.

When a connection is created, the station selects the first and last hop satellites from the constellation [StationUdlRouting]. The station then performs call admission control [StationQoSManager] to determine whether there are enough resources [StationResources] to support the connection. Then the connection request is passed to the first satellite which computes the route [IslRouting] between the first and last satellites. If there is such a route, all satellites on the path perform call admission control [SatelliteQoSManager], [SatelliteResources]. The same procedure takes place in the destination station. If it turns out that the connection can be routed through that path, resources are actually allocated ([StationQoSManager], [SatelliteQoSManager]).

In order to simulate connections coming from a call concentrator (aggregated phone calls), the number of channels of the connection is not fixed after a connection has been set up, but can change during the lifetime of the connection. For example, a concentrator may set up a single connection for all the phone calls it handles, and may simulate both new phone calls and old closed phone calls by varying the number of channels used by the single connection as set up at start time. In other words, a number of n phone calls from station i to station j is simulated by the generation, in station i , of a unique connection that requests n channels. The modification of a connection requirements is performed in a similar way to the connection setup procedure.

A *handover* (or hand-off) occurs when either a UDL connecting a satellite to a ground station is cut off, or when a beam change occurs (inside the same UDL), or when an ISL is cut off. All connections passing through the affected link must be appropriately processed (rerouted or torn down) [ConnectionChangeMonitor].

A *connection drop* occurs when an existing connection is forcibly torn down. It may happen either when there is a handover and the connection cannot be rerouted, or when higher priority traffic preempts all the resources used by a connection. A *partial drop* may also occur, when part of the resources of the connection is taken back by the network. A *call block* occurs when a new connection cannot be established. It may happen when there are no resources available in the network in order to support the new connection. At the current time, the limitations of the call connections are: only point-to-point connections are considered; a connection cannot be split on more than one path (however, forward and return channels are not necessarily on the same path); no rerouting of connections happens as a con-

sequence of growing or shrinking a connection (aggregate connections case); and no partial rerouting of connections is possible inside the constellation. Some or all of these limitations may be lifted at a future time.

4.3 Routing

Routing policies are one of the main aspects that will be studied using GaliLEO. End-to-end routing is split into UDL and ISL routing; additionally, UDL routing is split into Up-Link (UL) routing, that is, the process by which the source ground station selects the source satellite used to forward the packets of the connection, and Down-Link (DL) routing, that is, the process by which the destination ground station selects the destination satellite from which the packets of the connection will arrive. On the other hand, given a source satellite and a destination satellite, as provided by UDL routing, ISL routing computes the (or at least one) optimal path between these two satellites. ISL routing includes a signalling scheme [SatelliteLinkStateManager] to distribute and gather routing information [RoutingInformation] to/from the other satellites.

4.4 Fault management

The general reliability of a satellite must cope with the reliability of each element as well as the relationships among different failures. Trying to compute the reliability function of a system is thus quite complex and many simulations have to be used in order to have an estimate of the reliability function. GaliLEO will take advantage of the capabilities of Consim, which will generate failure events having proper nature and time distribution.

5 Some implementation aspects

This section will cover some implementation issues related to simulation performance. As it is often the case with broadband network simulations, the time needed to simulate a short period of time may be in the order of days; hence, the concern about performance enhancement. We will go briefly through considerations about simulating the network packet flows, distributing the simulation, programming optimisations and the selection of an appropriate development tool.

Simulating the actual packet flow in a network simulator provides valuable insight on the network behaviour. Without doing so, a satisfactory level of accuracy, especially when it comes to time relations of the various phenomena occurring in the network, cannot be achieved. Unfortunately, it also results in a heavy process (if not intractable)

for a simulator of LEO constellations because of the potential huge number of traffic sources, and because of the bandwidth ranges involved (up to hundreds of Mbit/s). As a result, simulating each packet individually is often only wishful thinking for realistic simulation scenarios. Two solutions are available to overcome this problem. The first solution consists in using mathematical tools (when it is possible) to model the average behaviour of the packets and deduce useful measures. For example, if the traffic is made of a number of constant bit rate sources, one can - given certain assumptions - model the cell arrival pattern in a switch using a $ND/D/1$ queue [11].

The second solution is to implement distributed or parallel simulation in order to increase the available processing power [8]. GaliLEO plans to support both solutions. The simulation engine of GaliLEO was designed in order to ease the transition to a distributed paradigm without compromising the existing architecture.

Implementing distributed simulation raises two issues. The first one is how to partition the processing space into parallel processing entities. The second one is implementation related and concerns the communication means that are used among processing entities. As far as GaliLEO is concerned, one possible partition is to distribute evenly the satellites and stations on the pool of available computers. In order to choose a suitable partition, each possible solution must be evaluated taking into account the amount of data that has to be exchanged between the various distributed entities, the balance of the computation load on the different entities, the time dependencies between the entities and the available resources.

Once a distribution scheme has been established, the communication means must be chosen. Commonly such a mechanism provides *remote function call* like services. Java supports a distribution paradigm through remote method invocation. In a medium term range, the simulation engine of GaliLEO and LeoSim have been scheduled to make use of the RMI or other facilities (such as MPI) provided by Java. A survey of the different solutions available as well as of their performance has still to be performed. Currently, the simulation engine optionally supports parallel event processing on multiprocessor computers.

In a sequential or distributed simulation environment, performance improvements can be achieved at the implementation or system level. Enhancements are either related to the algorithms and data structures or to the development tools. All algorithms and data structures which are likely to be used often during the simulation must be carefully chosen. The agenda in the simulation engine is an example. Since thousands, if not millions, of events will be generated, queued and processed during a simulation run, these operations have to be efficient.

Additional concerns are raised by the Java memory man-

agement system, which uses a garbage collector. Although garbage collection makes it convenient to write code less vulnerable to memory related bugs, this feature calls the developer for careful attention to the object lifetime. Among other things, favouring object reuse is crucial to minimising the number of allocations as well as the number of objects eligible for garbage collection. This problem has surprising ramifications: past experience has shown that LeoSim's execution speed was almost doubled by increasing the heap size, therefore reducing the number of times the garbage collector is invoked.

Java was initially born as a language for developing Internet applications and delivering them on different hardware architectures without re-compilation. Compiling a Java program produces an intermediate language called byte-code. When the Java program is executed, the Java Virtual Machine (JVM) interprets the byte-code. The JVM takes care of the mapping between the byte-code and the native host architecture. Nevertheless, Java can also be used to develop applications that do not require seamless cross platform execution. The byte-code interpretation phase is a drawback from a performance standpoint. The first solution is to translate directly a Java source in native machine code. The GNU Java compiler (*gcj*), while still in the development phase, provides such a facility. An intermediate solution is to use a JVM with a Just-in-Time (JIT) compiler that translates byte-code to native code upon class loading. Some measurements made with LeoSim showed that the increase in execution speed approaches 90%. These measurements were made using IBM's JDK under Linux. Other tests are carried out with Sun's HotSpot, and Symantec's JVM. The choice of Java may be arguable, especially from a performance standpoint. However, the technology related to Java compilation is rapidly evolving and Java compilers producing machine code rather than intermediate byte code are becoming more and more popular. Java lends itself naturally to build multi-threaded applications for running on multiple local or remote processors, it is very easily portable among systems with varying graphics user interfaces, and a Java program can be easily adapted to present its graphics interface on a web browser. Finally, Java is a safe language which permits fast prototyping and strong checking during compilation and execution.

6 GaliLEO project management

GaliLEO is a medium-sized project with several remotely located teams participating. An effective mean to exchange information is mandatory. Furthermore, as for all developments, a structured approach is required. GaliLEO's project life cycle is following a spiral approach based on a core simulator incrementally enhanced. The analysis and design rely heavily on diagrams as a universal communica-

tion medium. The diagrams follow the UML standard and internal guidelines.

The first stage of the project consisted in writing in plain text what were the objectives of GaliLEO and ordering them by priority. Then the interactions between the user and GaliLEO were roughly described (using interaction diagrams from UML [10]). Using these diagrams as a starting point as well as our previous experiences, the system was described in terms of collaborating objects (i.e. objects exchanging messages). Then, these objects were grouped in classes. The class descriptions consisted in Java stubs documented using the *javadoc* utility from the JDK. At this point, GaliLEO was successfully compiling, although no processing was being done. This approach made it possible to gradually fill the gaps (i.e. replacing stubs with method bodies) while being able to test almost immediately the resulting code.

All deliverables are available in HTML from a Web server. Similarly, the source code is stored in a Web CVS repository. The CVS repository takes care of the versioning and is a useful tool to determine the changes made by different parties across successive versions. Currently, the primary development and analysis platform is Linux. All applications that were used during the design (*tgif*) and the development (*JDK*, *CVS*, *cvsweb*) are available free of charge.

7 Project status

Work on the design of GaliLEO begun in September 1998. Until June 2000, six Short Term Scientific Missions were organised and funded under the COST 253 action budget. Two additional missions were funded by the CNUCE-CNR. GaliLEO progresses mostly during these missions since the people involved (approximately 2.5 persons from CNUCE-CNR (IT), ENST (FR) and Public University of Navarra (ES)) have their regular activities to carry on.

Currently, an initial version of GaliLEO is available with simple but operational components. Among them, a shortest path ISL routing algorithm, a resources management scheme using firm allocation, a call generator using Poisson arrivals and a handover resolution policy implementing complete rerouting.

As a simple demo, let us consider comparing static routing versus adaptive routing for ISL networks. *Static routing* consists of computing a route (i.e. a path between two satellites in the constellation) once for all, without taking into account the evolution of the network state (e.g. the formation of congestion points). *Adaptive routing*, on the other hand, integrates network state knowledge in the routing decision process information. However, an additional mechanism (the routing signalling) is required to distribute such network state information among the satellites of the

constellation.

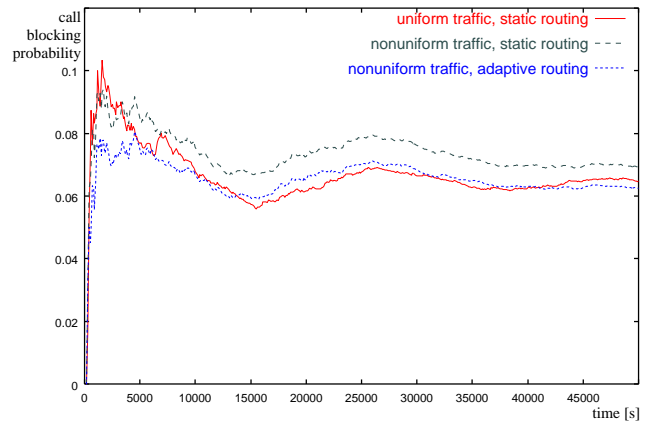


Figure 4. Connection blocking probabilities for three different scenarios.

The demo is composed of three steps which share some common parameters:

- There are twelve stations evenly distributed on the earth surface between latitudes 45° and -45° .
- Each station issues connection requests according to a deterministic scheme (i.e. the inter-arrival delay and the connection duration is fixed). For a given station, the destination station of the connection request is selected randomly among the other stations. Note: the generator seed is hardcoded in the input data, therefore for all simulations the sequence of selected destinations will be the same.
- Resources in the satellites, ISLs and stations are expressed as a number of acceptable connections. Each connection requires therefore one unit of resources. The constellation is based on Leonet, a 15 satellite LEO constellation with 3 orbits. ISLs accept 20 connections. Up/Down Links accept up to 80 connections. Satellites switch at wire speed.
- The performance measure is the connection blocking probability, that is, the ratio of the number of connections that could not be established to the total number of connection requests.
- The route taken by a connection is computed upon request, in the first satellite of the connection. The first and last satellites are chosen by the station initiating the connection request as those having maximum elevation.

The demo is made of three steps. In the first step, all stations issue the same traffic intensity (6 Erlangs). The routing algorithm is static. In the second one, a cluster of four

adjacent stations issue 10 Erlangs of traffic each, while the remaining stations issue 4 Erlangs each. The global traffic intensity is the same as in the previous case, but the traffic is not evenly distributed, so some satellites are congested, and the blocking probability increases.

In the third step, the traffic distribution is the same as in the second one, but the routing algorithm is adaptive. Every 15 seconds, each satellite broadcasts its available capacity to all the other satellites. When computing a route, the congested nodes and links are discarded. As a result, the connection blocking probability decreases again, because the routing algorithm is able to work around the congested nodes and links. On a Pentium I 200 MHz, each simulation requires approximatively 2 minutes.

8 Conclusions

Considering the questions still open in the field of LEO constellations, there is an urgent need for a simulation tool that would provide means to study these questions. GaliLEO is meant to be this tool and will, as a first step, be aimed at the study of constellation access techniques, routing algorithms, and fault management. GaliLEO is an ambitious project with many challenges which will provide in the end a valuable tool for the organisations involved in LEO research.

References

- [1] M. Annoni, S. Bizzarri, and F. Faggi. Performance evaluation of satellite constellations. the CONSIM(TM) simulator concept and architecture. In Springer-Verlag, editor, *Third European Workshop on mobile/personal Satcoms (EMPS'98)*, Venezia (IT), Sept. 1998.
- [2] R. Brown. Calendar queues: A fast O(1) priority queue implementation for the simulation event set problem. *Communications ACM*, 31:1220–1227, 1988.
- [3] M. Campione, K. Walrath, and A. Huml. *The Java Tutorial - Third Edition*. Addison Wesley, 2001.
- [4] N. Celandroni, E. Ferro, and F. Potorti. A simulation tool to validate and compare satellite TDMA access schemes. *Telecommunications Systems*, 12(1):21–37, 1999.
- [5] K. B. Erickson, R. E. Ladner, and A. LaMarca. Optimizing static calendar queues. In *35th IEEE Annual Symposium on Foundations of Computer Science*, pages 732–743, Santa Fe (US-NM), Nov. 1994.
- [6] E. Ferro. Proposal for a simulator architecture. Cost253 Temporary Document 10, CNUCE-CNR (IT), 1998.
- [7] L. Franck. *Routing for Inter-satellite Link Networks*. PhD thesis, Ecole Nationale Supérieure des Télécommunications, 2001.
- [8] R. M. Fujimoto. *Parallel and Distributed Simulation Systems*. Wiley, 2000.
- [9] MIL3. *Opnet tutorial manuals*.
- [10] A. Muller. *Instant UML*. Wrox, 1997.
- [11] J. Pitt and J. Schormans. *Introduction to ATM design and Performance*. Wiley, 1996.
- [12] C. H. Sauer, mac Nair, et al. Queuing network software for systems modelling. *Software – Practice and Experience*, 9(5), 1978.
- [13] C. H. Sauer, E. A. mac Nair, and J. F. Kurose. The research queuing package: past, present and future. In *Proceeding of the National Computer Conference AFEPS*, Arlington, Virginia (US), 1982.
- [14] K. S. Shanmugan. *Bones designer: introductory overview*. Comdisco systems.