

STORING AND EXCHANGING SIMULATION RESULTS IN TELECOMMUNICATIONS

Dragan Savić¹, Francesco Potorti², Francesco Furfari²,

Janez Bešter¹, Sašo Tomažič¹, Matevž Pustišek¹

¹University of Ljubljana, Faculty of Electrical Engineering,
1000 Ljubljana, Tržaška 25, Slovenia

²ISTI-CNR, Pisa, Italy

dragan.savic@fe.uni-lj.si (Dragan Savić)

Abstract

Though storing and exchanging simulation data is a rather simple task done by simulation practitioners, it is quite often a challenge as huge quantities of data are not uncommon, and conversion between different formats can be much time consuming. After examining some of the needs of the telecommunications simulation community, we describe the architecture of a working prototype – CostGlue – to be used as a general-purpose archiver and converter for large quantities of simulation data. The software architecture of the CostGlue tool is modular therefore allowing further development and contributions from other research sphere of activity. The core of the tool – CoreGlue – is responsible for communicating with the database. It acts as a unified interface for writing to the database and reading from it. Specific functions like import and export of data and different mathematical calculations are represented as a set of self-described modules, which are loaded as necessary. The graphic user interface is introduced as a web application for the simplicity of use and effective remote access to the application. The software package CostGlue is going to be released as free software with the possibility of further development.

Keywords: simulation, measurements, data, archiving, HDF

Presenting Author's biography

Dragan Savić received his B.S. degree in electrical engineering from Faculty of Electrical Engineering, University of Ljubljana, Slovenia in 2004. He is currently a Ph.D candidate in the Laboratory for communication devices. His research interests include simulation of communication networks and data mining in the field of Telecommunications. He is an IEEE member and a president of the IEEE student branch of Ljubljana.



1 Introduction

Simulation is quite common in the field of telecommunications because of the difficulties of many research challenges, which originate from the complexity, largeness and heterogeneity of the telecommunication networks that are being simulated. Because of this properties simulations represent indispensable, universal and cost efficient approach to the development of the new networks and consequentially also the development of new services and applications. Therefore there is a great variety of simulation programs in use. These programs – simulators – serve the purpose of inspecting different design approaches and studies of telecommunication networks.

Information about the behavior and operation of the modeled system is gathered from the simulation data, which represent the output of individual simulation programs. In order to achieve a good flow of the gained information among the researchers and research institutions it is necessary to take care of the exchange of the simulation data and post-processed data. Currently, simulation data exchange is hindered by the usage of different kinds of data formats in simulation software packages.

This topic has been addressed [1] in the framework of the European COST 285 Action "Modeling and Simulation Tools for Research in Emerging Multi-service Telecommunications"; a forum where researchers from all around Europe have periodically met to address issues related to the simulation of communications systems. The point made was that apparently no general purpose tools exist for exchanging big quantities of simulation data coming from different sources in different formats. Not only was the need for a common format for exchanging data highlighted, but also the need of feeding this data to different tools for post-processing, each requiring a different input format.

To better understand the scope of different requirements we define a reference model that encompasses the creation, flow and processing of data in the analysis of telecommunications systems. The main functional parts comprising such a model (simulators, data collectors, graphing tools, statistical tools) are covered by the many existing tools that are used by the research community; we focus on input data in the form of simulation results (e.g., ns-2 [5] traces), or measurements (e.g., Tcpdump traces). Raw data can be post-processed (e.g., packet delay) and the results stored separately from the raw data or complementing them, so that further analysis is possible based on both raw and post-processed data. Finally, the results can be exported in various widespread output formats, like tabular ASCII data or XML.

Data storage is based on the HDF5 [2] data format, which was selected after the analysis of the available options. HDF5 has been successfully applied in several scientific projects; it enables efficient data storage and lookup. Among the features most relevant to our purpose, it provides support for extremely large quantities of data, meta descriptors, and embedded compression. A set of programming libraries is available in C and Python as well as in numerous others languages, which simplifies software development.

We propose a prototype tool called CostGlue capable of packaging simulation output data into a common format, launching external post-processing modules and exporting data into various output formats. The basic mechanism provided by the HDF5 data format to store metadata is enriched, in order to support a more useful XML Schema language in defining the significant information about simulation data.

The next sections describe the proposed architecture of CostGlue into deeper detail, particularly its core - CoreGlue.

2 Formats for exchanging data

CostGlue acts as an archive for data generated by various different simulation programs and as a converter from several output formats to several input formats. Therefore, it is important to know which programs and formats are generally used by telecommunications systems practitioners. To this end we used the information obtained from the COST 285 participants, representatives of more than ten European nations, about the kind of tools they use for their simulation work. We learned that no single simulation tool has a dominant position, but that there is a great variety of tools in use.

A specific questionnaire directed to the above mentioned people yielded some interesting results:

- Apart from tabular data, other types of data organization, such as hierarchical or other elaborate structural formats are rarely used.
- Most of the time, data are used for statistics or graphing, other uses such as data mining are rare.
- The most common method for evaluating the statistical accuracy of simulation data is to use independent replications with the combination of different simulation times to assure long enough stationary intervals; equivalent correlation length (single simulation run) is rarely used; these methods are generally applied on a case-by-case basis, with the help of custom Bash, Perl or Python scripts.
- A single simulation run produces anywhere from 1 MB to 2 GB of data, and a simulation campaign requires from 1 to 100 runs, in the responses we got; measurements campaigns (as opposed to simulation)

required from 1 to 5 runs, each generating from 100 MB to 50 GB of data.

- Storage required varies from up to 1 GB for short-term storage to anywhere from 10 MB to 10 GB and more for long-term storage.
- Used metadata include type, date, parameter values and their description, version tracking, configurations, simulation scripts, and location.
- Metadata is stored in different locations: coded into directory and file names, in separate files, in different storage location (e.g., under root directory, shared directories) inside files, and in databases.

Among those who use a predefined output format, the most common appears to be the network simulator ns-2. Among the generic tools for mathematical computation and running simulations, Matlab appears to be used by many. A large part of the simulators is composed by standard scripting or programming languages and, in general, by ad hoc simulators. Concerning the tools used for post-processing or graphing, there is an even greater variety. These observations, while limited in scope, show that some sort of tabular ASCII format is of common use, and thus that being able to read and write ASCII tabular data is certainly a requisite for our proposed archiver and converter. Nevertheless, the variety of tools used also calls for a general way of reading and writing many formats: that is why we consider the modular architecture of CostGlue a necessary feature for the tool to be useful at all.

Another interesting point is that simulation data and measurement data have a lot in common, and a tool useful for one can be also useful for the other. However, measurements like simulation results are often output in particular formats, and an input converter is very frequently needed. An interesting feature that can be made part of CoreGlue is the ability to give a similar treatment to data coming either from measurement or from simulation of the same environment, and archive them in the same format. This is the reason why the first prototype of the CostGlue includes the ability to read ns-2 data and Tcpcdump data, store them into a common format and write in Nam, ns-2 and Tcpcdump formats. This capability makes it easy to use the many tools available that are able to analyze and graph data obtained by both ns-2 and Tcpcdump.

All the above discussion leads to a scenario where a simulation tool is run several times, each time producing tabular data, that is, data that can be conveniently stored into a two-dimensional structure having relatively few columns and a possibly large number of rows. What about data that cannot be naturally converted to a two-dimensional format? In this case, the inner structure of the archived data needs to be different. One of the challenges is to make the tool efficient in the most common case of collections

of tabular data, but still be useful in the case of non-tabular data.

3 The database and the database structure

A common file format solves several problems of simulation data exchange. Therefore we made a thorough analysis of different data formats and their corresponding libraries for data manipulation. Among many, we have focused on the following set of data formats: HDF4, HDF5, netCDF, ODB, FITS and OpenDX. Beside these, we also considered using plain text formats, XML and SQL databases. The results of the analysis makes it clear that the HDF5 file format is the most suitable for this task, since it meets all the requirements of data organization e.g., separation of raw data and metadata, and different requirements of contemporary computer system architectures, such as managing big quantities of data, offering a general data model, supporting complex data structures, portability among different computer platforms, parallel data access and processing, diversity of physical file storage media, and sustained development and maintenance.

Summarizing the above, most simulation data in the computer communications area are collections of tables of numeric data: each simulation run generates a table of data having few columns and a possible large number of rows. Each table is associated with certain parameters specific to a simulation run that generated it, and is uniquely identified by the values of those parameters. We are interested in defining a database structure that is able to efficiently accommodate this type of data. HDF5 meets these requirements.

HDF5 is data format with associated software library. The software library consists of two primary objects: dataset and group. A dataset represents a multidimensional array of data elements, which can hold different types of data. The data stored in datasets can be either homogeneous (only one data type used within a single dataset - simple dataset) or compound (different number of data types within one dataset - compound dataset). Since tabular data collected from certain simulators often contains data with different types (e.g., integer, float, char), we use compound datasets to accommodate the nature of simulation outputs. An HDF5 group is a structure containing zero or more HDF5 objects. By using two primary HDF5 objects, data can be organized hierarchically by means of a tree structure where an arbitrary number of HDF5 objects are derived from the main "root" group. Groups and datasets have a logical counterpart in directories and files in a hierarchical file system and, similarly to a file system, one can refer to an object in an HDF5 file by its full path name.

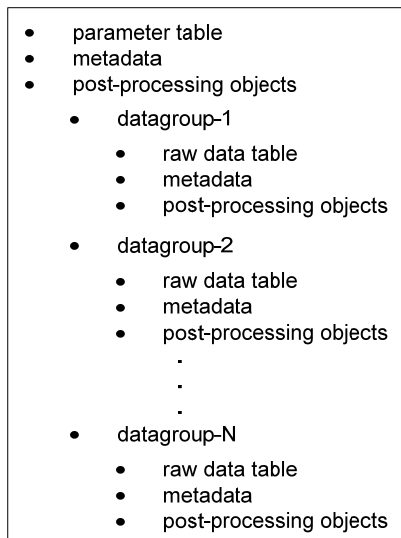


Fig. 1 General overview of the structure of the database

To meet the requirements of efficient data storage, especially those critical to management, understanding and reuse of scientific data, each HDF5 object may have associated metadata stored in the HDF5 file (referred as archive) in a simple attributes form. Attributes usually represent a small dataset connected to a certain group or a dataset. Their purpose is to describe the nature and/or the intended usage of the object they are attached to.

In the design of the database structure our goal was a flexible representation of the stored simulation data by using one multidimensional array, where a user can easily extract a desirable portion of the simulation data. Even though HDF5 supports multidimensional arrays, it is not efficient to store large amount of data in just one array. Furthermore, due to the HDF5 primary aspect of use, which is to have data organized hierarchically, storing everything inside a single multidimensional array would be a loss on the side of flexibility. We also introduced a parameter table which acts as an index. By this we are able to map the logical view of a multidimensional matrix into an HDF5 hierarchical structure.

Fig. 1 presents a general overview of the proposed database structure, where the parameter table represents the logical part and all other groups and datasets represent actual data, where the raw simulation data, metadata and post-processing data is stored. The whole database is treated as one archive containing a "root" group from which all other groups and datasets form a two-level tree.

An immediate extension to having tables of numeric data is having tables of fixed-length data, which can be flags, numbers or strings. The PyTables library allows efficient manipulation of 2-dimensional HDF5 compound datasets from Python referred to as tables from now on. Each table, together with metadata and post-processing data, is attached to a data group,

which usually holds the data produced during a single simulation run. Data groups are indexed by vectors of parameters representing individual simulation run. Thus the index is a 2-dimensional array, referred also as the parameter table, where the parameters relative to each data group are stored: each column corresponds to a different parameter, and each row contains the values of the parameter relative to a data group. Therefore an index is used as the data structure for accessing a data group, using an array of parameter values relative to that data group as the key. A table is attached to each data group, where each row is filled with the values of the fields, each field corresponding to a column of the table.




Fig. 2 Multidimensional view of the raw data

The overall structure is a collection of 2-dimensional tables indexed by arrays of P parameters, as described on the Fig. 2. This can be logically seen as a matrix with $N+2$ dimensions, where the first N dimensions are sparse and the last 2 dimensions ($N+1$ and $N+2$) are dense. We define the first N indices as parameters identifying a data group. As for the last two indices, the first (F) of them represents the field in the data group table, and the second index (R) the row number of the data group table.

Let us now describe how the results of an example real-world simulation can be stored in the described structure. We are simulating the behavior of packet switches in ns-2; each run is characterized by several parameters, such as architecture type, buffer size, number of I/O ports and traffic load. Each simulation run differs in at least one of these parameters. When storing the results of one simulation run in the archive, we populate a new row in the indexing table: values in each row uniquely identify a simulation run. Each row includes the full path to a data group, containing the table with the results of the simulation run, metadata and processing data, as shown in Fig. 1. Metadata stores information about the type of scripts used to generate that simulation run, the type of network topology, traffic patterns, etc. Post-processing data include packet loss probability, maximum, minimum and average packet delay. Metadata and processing data are also associated with the whole archive, and contain information relative to the whole set of simulation runs.

3.1 Metadata and processing data

In recent years increasing attention has been devoted to metadata for every application domain. The XML Schema language [7] provides a means for defining the structure, contents and semantics of an XML document and it is widely used to collect data about data, that is, metadata. The HDF5 data format allows

metadata to be associated with every object by using a series of predefined attributes in the form of name=value pairs. This mechanism is too simple for our requirements, so we decided to use it just to replicate some relevant piece of information, in order to increase the robustness of the archiver system. Consequently, in order to insert metadata, we defined an XML Schema whose document instances can be saved together with the simulation data, as reported in Fig. 1. Metadata can be associated to every data group; metadata referring to the archive as a whole are saved together with the indexing table, while metadata for a single simulation run are saved in the related data group. Metadata can make reference to any kind of additional data, which are labeled in Fig. 1 as post-processing objects; examples are statistics on the raw data, charts, images, and any type of data that are produced from or relevant to the raw data.

The CostGlue metadata XML Schema is derived from the Information Model defined in the OAIS reference model [8] and from the Scientific Data Model (CSMD) [9]. The Open Archival Information System (OAIS) is a technical recommendation to provide permanent or indefinite long-term, preservation of digital information. Of particular interest is the OAIS Information Model. The preserved information is called content information and is accompanied by a complete set of metadata as depicted in Fig. 3, where the content data object is the actual preserved data (a simulation run or measurement trace in our case) and the representation information consists of the information that is required to render and interpret the object. This might include the specification of the data format as well as the software needed to access the data. The reference information, used to provide assigned identifiers to the content information, allows to univocally referring the object outside of the system. Typical identifiers could be UUIDs or terms of taxonomy. The context information documents the relationships between the content information and its environment, including why the object was created and how it is related to other objects. The provenance information specifies the origin of the content information, any changes that may have taken place and so on. Finally the fixity information provides data integrity checks to verify that the object has not been altered in an undocumented way.




Fig. 3 Archive package as defined in the OAIS information model

Our XML Schema is based on the OAIS information model with a structure specialized for describing a

collection of simulation runs. Indeed, the adoption of an XML schema could facilitate further aggregation of telecommunication archives in catalogues to be published in repositories for the scientific communities [4, 10]. For this reason we have adopted in our schema the taxonomy used in the CSMD model depicted in Fig. 4, where the boxes labeled experiment, simulation and measurement contain the raw data we aim at archiving. The objective of the CSMD model is to aid interoperability of scientific information systems among research organizations. The information related to the project or investigation to which the simulation data refer are optional but, if provided, they are used in both the reference information and in the context information metadata.

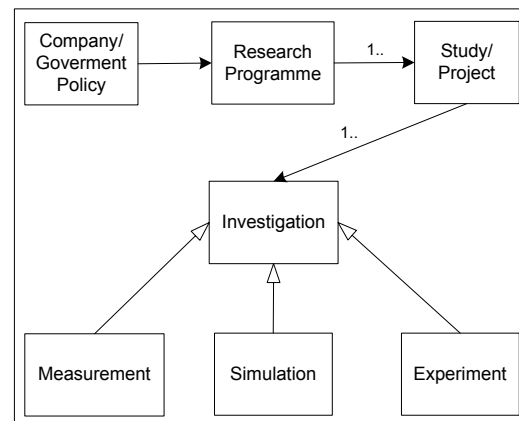


Fig. 4 Main classes of the CSMD model version 2

4 CostGlue architecture

The CoreGlue manages the index and the database structure, including the tables. Modules are responsible for metadata and post-process contents, both for the single data groups and for the whole archive. The CoreGlue and the modules together constitute the whole application, which is named CostGlue.

The CostGlue is written in Python. This language was chosen because of anecdotal evidence of its efficiency both in memory usage and processing power and for its programmability ease, due to automatic garbage collection and many native functions and types. Portability among operating systems is excellent and library availability for many tasks, especially mathematical ones, is rich. With respect to its main competitor, Java, Python has a generally smaller memory footprint and since it is completely free, does not suffer from being controlled by a single entity.

The architecture of the CostGlue application, depicted in Fig. 5, consists of a core, called CoreGlue, and several specialized modules. The core takes care of reading and writing to the HDF5 archive and of the dynamic loading of modules, while the modules are devoted to specialized tasks. Modules interact with the core through an API. Examples of tasks a module can perform are:

- Data import/export to/from the archive.
- Statistical computations over data stored in the archive.
- Data extraction from the archive with complex filters.
- Transformations over the data contained in the archive.
- Graphical output creating from the data in archive.
- Provision of a generic graphical user interface (GUI) for exploring the data in the archive, doing simple import/export or computations and running other available modules.
- Provision of a module-specific GUI.

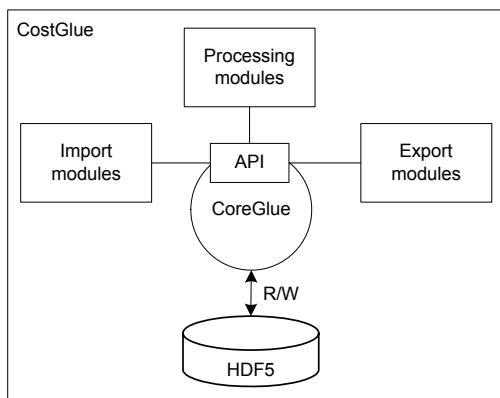


Fig. 5 CostGlue architecture

5 Overview of the module API

Python, like many modern languages, supports dynamic module loading. We exploit this possibility by providing a module library with a limited number of modules and leaving it open for other parties to write new modules. Modules are Python self-describing programs, which reside in a fixed place. The CoreGlue can look for available modules and query them one by one in order to get to know their capabilities; this can be used for instance for building a menu for a GUI (graphical user interface). Modules are able to describe the parameters they need and the type of output they produce. The needed parameters can then be provided to the module by the CoreGlue with or without input from the user; in the case where input is required, CoreGlue checks the provided parameters for consistency. A GUI can also use this information and present the choices to the user. Modules interact with the CoreGlue through a well-defined API which contains all necessary classes and methods for interacting with the database. The methods allow a module to manipulate the data group index in order to add or remove data groups, and to manipulate the data groups, in order to add or remove data, metadata and post-processed data.

The API also includes methods for accessing data with selectors written in index notation, which is widely

used in matrix computations programs such as Matlab or Octave [3]. In this notation each of the P+2 indices can be "1", indicating the smallest index; "end", indicating the highest one; "n:m", indicating all the indices between n and m included; ":", indicating the whole range from smallest to highest, "n:s:m", indicating the range from n to m in steps of s; an array like "[1 5 6 8]" indicating the selected indices. Using the Matlab index notation one can take complex orthogonal slices of the multidimensional matrix composed of all the data in the database; in fact, the database structure can be seen as a sparse matrix with P+2 dimensions, and being able to take a slice of this matrix could prove a powerful feature of CostGlue.

The CoreGlue contains three different queues: command, result and report queue, which are used for exchanging different types of messages with modules. When using the CostGlue the first step is to call the CoreGlue together with a module name and its optional arguments. A given module can work either alone, e.g., by performing a batch job, or it can depend on other modules, e.g., a command line interface can call a specialized module to do some work. Modules can also differ in the way they are built. Simple ones, after sending the command to the CoreGlue, wait for the results and are therefore not able to perform any additional tasks. Alternatively, complex modules run asynchronously with respect to the core, so they can check periodically the result queue for results and in the meantime inform the user about the current work progress by reading report messages from the report queue. Modules working as described are run as separate threads inside the core.

5.1 The command line and the HTML GUI modules

When CoreGlue is invoked by a command line, its first argument is the module name, followed by the parameters needed by the module. When invoked with the argument `html`, the CoreGlue acts as an HTTP server, providing a graphical user interface. Through this interface, the user can look at the list of available modules and, for each of them, look at a description, at the input they require and at the output they provide. The CoreGlue provides an interface for the input of module parameters, complete with checking, thanks to the information that it reads from the modules themselves. In the simplest case, this is analogous to calling the module on the CoreGlue command line, but more convenient for interactive use. A module can also provide a graphical interface or graphical output by itself.

Currently an interactive command line module is implemented. This module is used for debugging of the CoreGlue and modules. Additionally, it can be used to import and export tabular data. This requires the user to specify options and arguments; a specialized module can automatically recognize the

type of a trace and automatically provide naming, and possibly, filtering.

5.2 Modules for antenna measurements and simulations

We have designed two different import modules together with researchers from SSR (Signals, Systems and Radio) group at the University of Madrid. The first module deals with antenna measurements. Measurements are taken in the special antenna chamber. There are three types of different kinds of measurements depending on the way in which the probe or the testing antenna are shifted: spherical (roll over azimuth – Θ and Φ), cylindrical or planar. Each measurement is performed on a single frequency and the results are saved in two separate files. The first file contains the information about the organization of the raw data, which also represents the second file. The second file is in a binary form. To be stored into HDF5 database (in form of tables) it needs to be converted into ASCII format. After the import action of all the performed measurements is finished we can start using the SSR's proprietary visualization module (written in Matlab) for visualizing the antenna's directivity diagrams. Each measurement is selected by a set of parameters from the visualization module. In this case the names of parameters are: client name, project name, antenna type and working frequency.

The second specialized module is for dealing with antenna simulation results. The results are a direct output of the GRASP8 antenna simulator. The whole process (import, export and visualization) is quite similar to the one described above. The only difference is that here the simulation results are already in ASCII format and there is no additional conversion. However because of the file structure we need to first extract the information from the header lines before the actual results are imported in HDF5 database. A very useful and practical feature is the ability of comparing the measurement results with simulated ones. Therefore the visualization module supports this activity by plotting diagrams from both types of antenna data.

6 Conclusions

The purpose of the conversion and storage tool described in this paper is to facilitate the exchange and management of simulation data among researchers, and to ease the task of using different simulation, measurement, data processing and visualization tools, all having different input and output data formats. A prototype with a debugging interface is already working. As soon as the prototype passes its testing stage, it will hopefully be enhanced by other simulation practitioners thanks to its modular structure. Possible extension to the presented architecture includes being able to accommodate non-tabular data.

We believe that software developed as part of research activity should be released with a free software license, because research results should be made available for use by anyone, for any purpose and be freely modifiable, in order to further knowledge and usefulness [6]. The choice of license will be among the MIT X license, the GNU LGPL, and the GNU GPL licenses which we think best serve the purpose of free research software.

7 References

- [1] BRAGG, A. Observations and thoughts on the possible approaches for addressing the tasks specified in the COST 285 work-plan. COST 285 temporary document TD/285/03/15, CNUCE-CNR (IT), 2004.
- [2] FOLK, M., MCGRATH, R., AND YEAGER, N. HDF: an update and future directions. In International Geoscience and Remote Sensing Symposium (IGARSS'99) (1999), IEEE, Ed., vol. 1, pp. 273–275.
- [3] GOLUB, AND LOAN, V. Matrix Computations, 3 ed. The Johns Hopkins University Press, 1996.
- [4] KOTZ, D., AND HENDERSON, T. CRAWDAD: A Community Resource for Archiving Wireless Data at Dartmouth. IEEE Pervasive Computing 4, 4 (Oct. 2005), 12–14.
- [5] MCCANNE, S., AND FLOYD, S. The network simulator - ns-2. University of Berkley, Oct. 2005.
- [6] POTORTÌ, F. Free software and research. In proceedings of the International Conference on Open Source Systems (OSS) (July 2005), M. Scotto and G. Succi, Eds., ECIG Edizioni Culturali Internazionali Genova, pp. 270–271. Short paper.
- [7] <http://www.w3.org/XML/Schema>
- [8] ISO 14721:2003, Blue Book. Issue 1. CCSDC 650.0-r-2: Reference Model for an Open Archival Information System
- [9] Shoaib Sufi, Brian Mathews, CCLRC scientific metadata model: Version 2, CCLRC technical report DL-TR-2004-001, September 2004.
- [10] DRIVER (<http://www.driver-repository.eu/>)