Universidad Carlos III de Madrid
Computer Science and Engineering Department

Computer Architecture and Technology Area

ComBoS - User Manual

Saúl Alonso Monsalve
Félix García Carballeira
Alejandro Calderón Mateos

Leganés, Madrid, Spain
June 2016

Contents

1 Introduction

This document presents a detailed user manual of Complete BOINC Simulator (ComBoS). First we indicate the basic requirements to deploy the application and a detailed tutorial for the installation of SimGrid. Finally, we present an example of the simulator usage.

2 Basic Requirements

The technical specifications recommended for the final user to obtain the best experience from the application are:

- **Operating System**: Ubuntu 14.04.4 LTS (Linux distribution) or higher.
- **Processor**: Intel(R) Core(TM) i7 CPU 920 @2.67GHz or higher.
- **Random-Access Memory (RAM)**: 8 GB or higher.
- **Storage**: 1 GB of free space in the Hard Disk Drive.
- **Network**: Internet connection is not required.
- **Software**: The following software must be installed in order to run the application:
    1. GCC (GNU Compiler) 5.1 or higher.
    2. SimGrid toolkit 3.10 or higher.

3 SimGrid Installation

We will present a tutorial for the installation of the SimGrid toolkit (version 3.10). First, you have to download the official binary package from the *download page* (`http://simgrid.gforge.inria.fr/download.php`). In this case you will download the file *SimGrid-3.10.tar.gz*.

Then, you have to recompile de archive. This should be done in a few lines:

```
$ tar xf SimGrid-3.10.tar.gz
$ cd SimGrid-3.10
$ cmake -DCMAKE_INSTALL_PREFIX=/opt/simgrid $HOME
```

```
$ make
$ make install
```

After following these steps, you will have the SimGrid toolkit installed in your computer.

4 Usage Example

In order to use ComBoS, you must download the corresponding files from the following website: `https://www.arcos.inf.uc3m.es/~combos/`. After unzipping the downloaded file, the unzipped files will follow the folder structure presented in Figure 4.1. To perform simulations using ComBoS, it is necessary to model the platform to be simulated. Once you know the environment to simulate, you must specify all simulation parameters in the parameters Extensible Markup Language (XML) file.

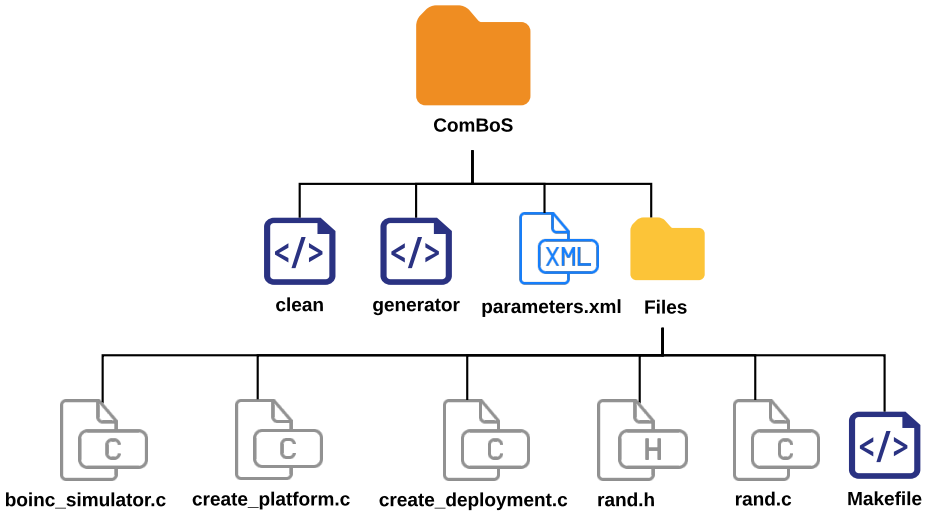

Figure 4.1: *Folder structure of the simulator.*

Figure 4.2 shows an example of a potential simulation that can be carried out by ComBoS. The figure shows a simplified platform with two Berkeley Open Infrastructure for Network Computing (BOINC) projects and 350,000 clients. The first project is represented by two scheduling servers (SS0 and SS1) and two data servers (DS0 and DS1). The second project consists of a single scheduling server (SS2) and three data servers (DS2, DS3 and DS4). Clients are grouped into three sets. The first group (G0) consists of 100,000 hosts and has a route to the first project. The second group (G1), has 200,000 hosts and a route to both projects. The third

group (G2) consists of 50,000 computers and has route to the second project. The rest of the figure shows the links among the elements of the environment (from L0 to L7). In each of the links, latency and bandwidth are indicated.

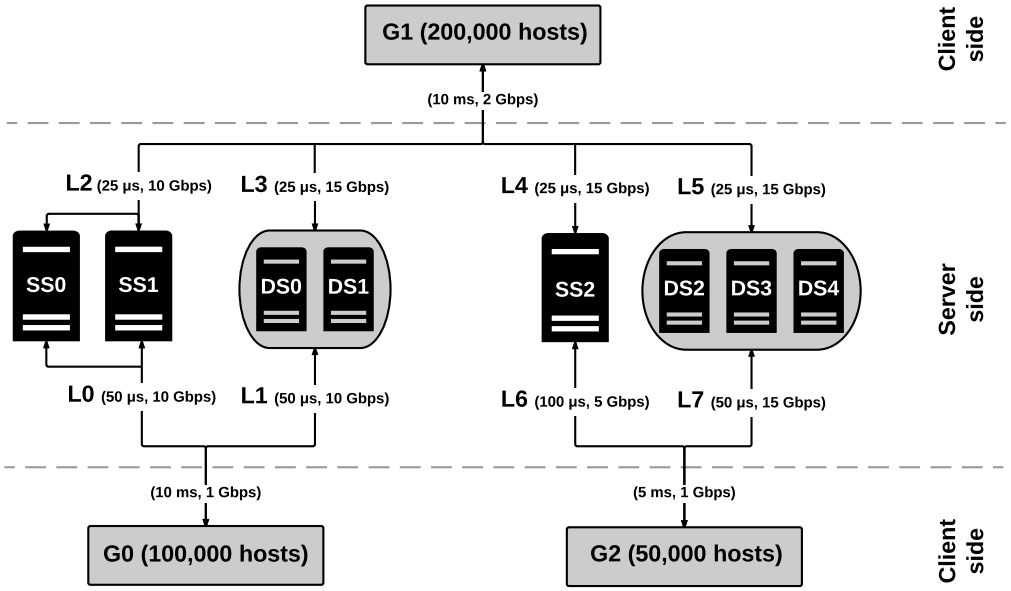

Figure 4.2: *Simulator platform example.*

To create a simulation, ComBoS requires to specify all the parameters described in Tables 1 and 2 in the parameters.xml file (see Listing 1). Users can define the power and availability of the volunteer hosts via either a traces file or distribution functions. For example, in the case of the SETI@home project, we have analyzed the 3,900,000 hosts that participate in this project. The CPU performance of the hosts can be modeled according to an exponential function, as shown Figure 4.3, which has a mean of 5.871 GigaFLOPS per host.

Our software first processes the XML file, creating the necessary deployment and platform files for subsequent simulations. In ComBoS, all this is transparent to the user. The user only has to specify all the parameters of the simulation in the XML file (Listing 1), and run the generator script using the following command:

```
$ ./generator
```

The above command generates the platform and deployment files. The platform file contemplates all the necessary elements in the simulation: hosts, clusters, links, etc. The deployment

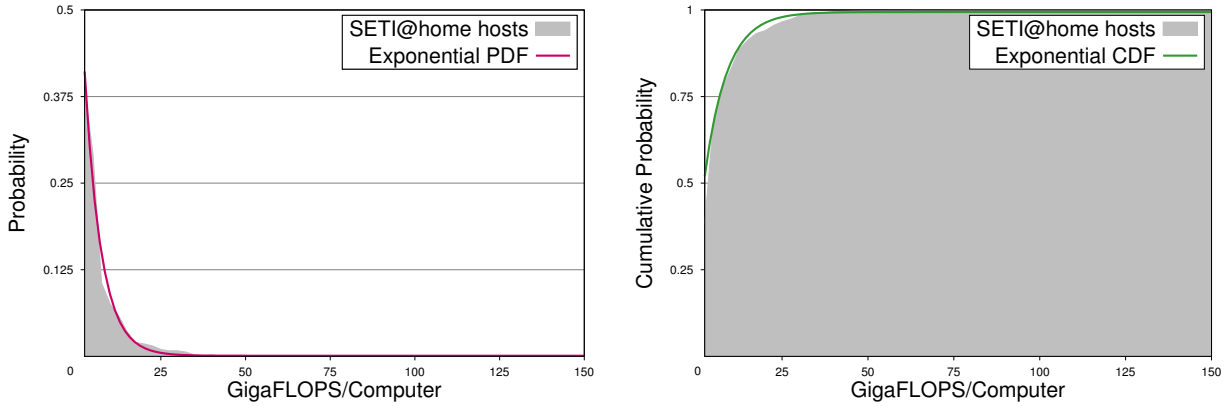

(a) *Probability density function of SETI@home hosts power.*

(b) *Cumulative distribution function of SETI@home hosts power.*

Figure 4.3: *CPU performance modeling for SETI@home hosts.*

file indicates the processes that should be created during the simulation. In addition, the generator script also compiles all source files needed for the simulation and generates the executable file. Finally, to run the simulation you just need to run the execution script:

```
$ ./execute
```

The execution results are composed by multiple statistical results (see Listings 2): the execution time, the memory usage of the simulator, the load of the scheduling and data servers, the total number of work requests received in the scheduling servers, the job statistics (number of jobs created, sent, received, analyzed, success, fail, too late, etc), the credit granted to the clients, the number of FLOPS, the average power of the volunteer nodes, and the percentage of time the volunteer nodes were available during the simulation.

Parameter	Description
name	Project name.
nscheduling_servers	Number of scheduling servers of the project.
ndata_servers	Number of data servers. If zero, scheduling servers also operate as data servers.
server_pw	CPU power of each server, in FLOPS.
disk_bw	Hard disk drive performance for each server, in bytes/s.
ifgl_percentage	Percentage of input files that must be generated locally on the client.
ifcd_percentage	Percentage of times a client must download new input files (due to locality scheduling).
input_file_size	Average amount of data that clients should download per workunit, in bytes.
output_file_size	Average amount of data that clients should upload per workunit, in bytes.
replication	Number of replicas of each file in the system.
task_fpops	Average task duration, in number of floating point operations needed to compute each task.
delay_bound	The time by which the result must be completed by the clients.
min_quorum	Minimum number of successful results required for the validator. If a strict majority agree, a consensus has been reached and the workunit is considered correct (there is a canonical result).
target_nresults	Number of results to create initially per workunit.
max_error_results	If the number of client error results exceed this, the workunit is declared to have an error.
max_total_results	If the number of results for this workunit exceeds this, the workunit is declared to be in error.
max_success_results	If the number of success results for this workunit exceeds this, and a consensus has not been reached, the workunit is declared to be in error.
success_percentage	Percentage of success results (when completed).
canonical_percentage	Percentage of success results that make up a consensus.

Table 1: *Project parameters.*

Parameter	Description
nclients	Number of VN of the group.
connection_interval	The typical time between periods of network activity.
scheduling_interval	The "time slice" of the BOINC client CPU scheduler (the default is one hour).
gbw	Bandwidth between each VN and the network backbone.
glatency	Latency between each VN and the network backbone.
traces_file	File with the VN power and availability traces (optional).
pv_distri	VN power fit distribution: Weibull, Gamma, Lognormal, Normal, Hyperexponential, Exponential (in case there is not a traces file).
max_power	Maximum power a VN might have using a random distribution.
min_power	Minimum power a VN might have using a random distribution.
av_distri	VN availability fit distribution: Weibull, Gamma, Lognormal, Normal, Hyperexponential, Exponential (in case there is not a traces file).
nv_distri	VN non-availability fit distribution: Weibull, Gamma, Lognormal, Normal, Hyperexponential, Exponential (in case there is not a traces file).
att_projs	Number of projects attached for each VN.
For each project:	
priority	Priority of the project, used by the client scheduler.
lsbw	Network bandwidth between the VN group and the scheduling servers of the project.
lslatency	Network latency between the VN group and the scheduling servers of the project.
ldbw	Network bandwidth between the VN group and the data servers of the project.
ldlatency	Network latency between the VN group and the data servers of the project.

Table 2: *VN group parameters.*

```
<simulation_time>100</simulation_time>

<!-- Server side -->
<server_side>
    <n_projects>2</n_projects>
    <sproject>
        <snumber>0</snumber>
        <name>PROJECT1</name>
        <nscheduling_servers>2</nscheduling_servers>
        <ndata_servers>2</ndata_servers>
        <server_pw>12000000000.0</server_pw>
        <disk_bw>80000000</disk_bw>
        <ifgl_percentage>100<ifgl_percentage>
        <ifcd_percentage>100<ifcd_percentage>
        <input_file_size>368640</input_file_size>
        <task_fpops>7560000000000</task_fpops>
        <output_file_size>65536</output_file_size>
        <min_quorum>2</min_quorum>
        <target_nresults>2</target_nresults>
        <max_error_results>2</max_error_results>
        <max_total_results>4</max_total_results>
        <max_success_results>3</max_success_results>
        <delay_bound>100000000</delay_bound>
        <success_percentage>95</success_percentage>
        <canonical_percentage>95</canonical_percentage>
        <replication>2</replication>
    <sproject/>
    <sproject>
        <snumber>1</snumber>
        <name>PROJECT2</name>
        <nscheduling_servers>1</nscheduling_servers>
        <ndata_servers>3</ndata_servers>
        <server_pw>12000000000.0</server_pw>
        <disk_bw>60000000</disk_bw>
        <ifgl_percentage>100<ifgl_percentage>
        <ifcd_percentage>100<ifcd_percentage>
        <input_file_size>52428800</input_file_size>
        <task_fpops>20800000000000</task_fpops>
        <output_file_size>16777216</output_file_size>
```

```
        <min_quorum>2</min_quorum>
        <target_nresults>2</target_nresults>
        <max_error_results>2</max_error_results>
        <max_total_results>4</max_total_results>
        <max_success_results>3</max_success_results>
        <delay_bound>100000000</delay_bound>
        <success_percentage>95</success_percentage>
        <canonical_percentage>95</canonical_percentage>
        <replication>2</replication>
    <sproject/>
</server_side>

<!-- Client side -->
<client_side>
    <n_groups>3</n_groups>
    <group>
        <n_clients>100000</n_clients>
        <connection_interval>1</connection_interval>
        <scheduling_interval>3600</connection_interval>
        <gbw>1Gbps</gbw>
        <glatency>10ms</glatency>
        <traces_file>NULL</traces_file>
        <max_speed>117.71</max_speed>
        <min_speed>0.07</min_speed>
        <pv_distri>5</pv_distri>
        <pa_param>0.1734</pa_param>
        <pb_param>-1</pb_param>
        <av_distri>0</av_distri>
        <aa_param>0.393</aa_param>
        <ab_param>2.964</ab_param>
        <nv_distri>2</nv_distri>
        <na_param>2.844</na_param>
        <nb_param>-0.586</nb_param>
        <att_projs>1</att_projs>
        <gproject>
            <pnumber>0</pnumber>
            <priority>1</priority>
            <lsbw>10Gbps</lsbw>
            <lslatency>50us</lslatency>
```

```
79          <ldbw>10Gbps</ldbw>
80          <ldlatency>50us</latency>
81      </gproject>
82  </group>
83  <group>
84      <n_clients>200000</n_clients>
85      <connection_interval>1</connection_interval>
86      <scheduling_interval>3600</connection_interval>
87      <gbw>2Gbps</gbw>
88      <glatency>10ms</glatency>
89      <traces_file>NULL</traces_file>
90      <max_speed>117.71</max_speed>
91      <min_speed>0.07</min_speed>
92      <pv_distri>5</pv_distri>
93      <pa_param>0.1734</pa_param>
94      <pb_param>−1</pb_param>
95      <av_distri>0</av_distri>
96      <aa_param>0.393</aa_param>
97      <ab_param>2.964</ab_param>
98      <nv_distri>2</nv_distri>
99      <na_param>2.844</na_param>
100     <nb_param>−0.586</nb_param>
101     <att_projs>2</att_projs>
102     <gproject>
103         <pnumber>0</pnumber>
104         <priority>1</priority>
105         <lsbw>10Gbps</lsbw>
106         <lslatency>25us</lslatency>
107         <ldbw>15Gbps</ldbw>
108         <ldlatency>25us</latency>
109     </gproject>
110     <gproject>
111         <pnumber>1</pnumber>
112         <priority>1</priority>
113         <lsbw>15Gbps</lsbw>
114         <lslatency>25us</lslatency>
115         <ldbw>15Gbps</ldbw>
116         <ldlatency>25us</latency>
117     </gproject>
```

```
118     </group>
119     <group>
120         <n_clients>50000</n_clients>
121         <connection_interval>1</connection_interval>
122         <scheduling_interval>3600</connection_interval>
123         <gbw>1Gbps</gbw>
124         <glatency>5ms</glatency>
125         <traces_file>NULL</traces_file>
126         <max_speed>117.71</max_speed>
127         <min_speed>0.07</min_speed>
128         <pv_distri>5</pv_distri>
129         <pa_param>0.1734</pa_param>
130         <pb_param>−1</pb_param>
131         <av_distri>0</av_distri>
132         <aa_param>0.393</aa_param>
133         <ab_param>2.964</ab_param>
134         <nv_distri>2</nv_distri>
135         <na_param>2.844</na_param>
136         <nb_param>−0.586</nb_param>
137         <att_projs>1</att_projs>
138         <gproject>
139             <pnumber>1</pnumber>
140             <priority>1</priority>
141             <lsbw>5Gbps</lsbw>
142             <lslatency>100us</lslatency>
143             <ldbw>15Gbps</ldbw>
144             <ldlatency>50us</latency>
145         </gproject>
146     </group>
147 </client_side>
```

Listing 1: *parameters.xml file filled with the parameters of the example.*

```
Memory usage: 11,570,812 KB

Total number of clients: 350,000

#################### PROJECT1 ####################

Simulation ends in 100 h (360,000 sec)

Scheduling server 0: Busy: 13.4%
Scheduling server 1: Busy: 13.4%
Data server 0:  Busy: 12.2%
Data server 1:  Busy: 12.2%

 Number of clients: 300,000
 Messages received:   32,227,795 (work requests received + results received)
 Work requests received:  16,179,688
 Results created:   16,179,689 (100.0%)
 Results sent:   16,179,688 (100.0%)
 Results received:   16,048,107 (99.2%)
 Results analyzed:   16,048,107 (100.0%)
 Results success:   15,245,637 (95.0%)
 Results failed:   802,470 (5.0%)
 Results too late:   0 (0.0%)
 Results valid:   13,616,636 (84.8%)
 Workunits total:   7,716,639
 Workunits completed:   6,863,142 (88.9%)
 Workunits not completed:  853,497 (11.1%)
 Workunits valid:   6,808,318 (88.2%)
 Workunits error:   54,824 (0.7%)
 Throughput:   89.5 mens/s
 Credit granted:   231,482,812 credits
 FLOPS average:   285,949 GFLOPS

#################### PROJECT2 ####################

Simulation ends in 100 h (360,000 sec)

Scheduling server 0: Busy: 1.7%
Data server 0:  Busy: 100.0%
```

```
40 Data server 1: Busy: 100.0%
41 Data server 2: Busy: 100.0%
42
43  Number of clients: 250,000
44  Messages received:  2,070,120 (work requests received + results received)
45  Work requests received:  1,161,080
46  Results created:  1,161,081 (100.0%)
47  Results sent:  1,161,080 (100.0%)
48  Results received:  909,040 (78.3%)
49  Results analyzed:  909,040 (100.0%)
50  Results success:  863,710 (95.0%)
51  Results failed:  45,330 (5.0%)
52  Results too late:  0 (0.0%)
53  Results valid:  743,688 (81.8%)
54  Workunits total:  559,384
55  Workunits completed:  374,713 (67.0%)
56  Workunits not completed:  184,671 (33.0%)
57  Workunits valid:  371,844 (66.5%)
58  Workunits error:  2,869 (0.5%)
59  Throughput:  5.8 mens/s
60  Credit granted:  35,697,024 credits
61  FLOPS average:  42,968 GFLOPS
62
63 Group 0. Average speed: 6.704465 GFLOPS. Available: 62.1% Not available 37.9%
64 Group 1. Average speed: 5.455870 GFLOPS. Available: 55.3% Not available 44.7%
65 Group 2. Average speed: 6.110686 GFLOPS. Available: 57.0% Not available 43.0%
66
67 Clients. Average speed: 5.906157 GFLOPS. Available: 57.5% Not available 42.5%
68
69 Execution time:
70 0 days 7 hours 51 min 49 s
```

Listing 2: *Simulation execution results.*

Universidad Carlos III de Madrid
Computer Science and Engineering Department

Computer Architecture and Technology Area

ComBoS - User Manual

Saúl Alonso Monsalve
Félix García Carballeira
Alejandro Calderón Mateos

Leganés, Madrid, Spain
June 2016