

libcppsim: A SIMULA-LIKE, PORTABLE PROCESS-ORIENTED SIMULATION LIBRARY IN C++

Moreno Marzolla

Dipartimento di Informatica
Università Ca' Foscari di Venezia
via Torino 155, 30172 Mestre (VE), Italy
e-mail: marzolla@dsi.unive.it

KEYWORDS Process-Oriented Simulation, Simulation Library, C++.

ABSTRACT

In this paper we describe the design and implementation of `libcppsim`, a general-purpose, process-oriented simulation library written in C++. `libcppsim` provides a set of classes for implementing simulation processes, scheduling primitives, random variate generation and output data analysis functions. The main simulation entity provided by the library is the simulation process; the basic process scheduling primitives are modeled upon those provided by SIMULA's simulation class. The modular object-oriented design of `libcppsim` allows users to extend its functionalities with minimal effort. In order to improve efficiency, simulation processes are not implemented as operating system threads; instead, we use coroutine objects which implements a cooperative quasi-parallel process environment. A portable implementation of coroutines is provided in order to allow `libcppsim` to be used on different platforms.

1 INTRODUCTION

Simulation is a general modeling technique which can be used for analyzing systems which can not be described and evaluated analytically. Many tools and languages have been developed in order to facilitate the implementation of simulation models. SIMULA [Dahl and Nygaard, 1966] has been one of the first successful general-purpose programming languages providing special features targeted at simulation implementation purposes. SIMULA provides a set of basic facilities for describing a simulation model as a set of cooperating simulation processes executing in a quasi-parallel system. These facilities allow the description of many typical simulation models, and could also be used as building blocks for more sophisticated applications; DEMOS [Birtwistle, 1979] was one of the alternative simulation systems developed in SIMULA.

Traditional general-purpose programming languages can be used for implementing simulation programs as well. Support for simulation facilities is usually pro-

vided as libraries which can be linked to user's applications. Examples of simulation libraries based on general-purpose languages include CSIM19 [Schwetman, 2001] and C-Sim [Hlavička et al., 1999] (based on the C language), C++Sim [Little and McCue, 1993], OMNET++ [Varga, 2001] and DESP-C++ [Darmont, 2000] (based on the C++ language) and JavaSim [Little, Little] (written in Java). More sophisticated and user-friendly simulation tools have also been developed to make the implementation of simulation models easier for non-programmers. Such simulation environments are usually limited to particular models, such as those based on Queuing Networks or Petri Nets; they provide users with visual editors for graphically building the model.

Simulation languages and libraries can be based on an event-oriented or process-oriented paradigm. Event-oriented simulation models are described in term of the events which can change the system's state. A process-oriented simulation model is represented as a collection of concurrently executing simulation processes, each one with its own thread of execution. A simulation process is made of two parts: a set of local variables, and a sequence of actions to execute.

It should be noted that many existing simulation languages and tools suffer either a long learning time, poor performances, or can only be used for special-purpose models as they lack generality. Simulation libraries based on the event-oriented paradigm, while usually easier to learn, are not suited for models with a high number of different event types, as the resulting simulation programs would be very difficult to understand and debug. On the other hand, implementing the process-oriented simulation paradigm on the top of a conventional, general-purpose programming language can be difficult due to limitations (the main being the lack of support for coroutines) of the host language. Nevertheless, using a standard programming language is desirable because modelers are more likely to be already acquainted with it and do not require to learn some new idiom.

In this paper we describe `libcppsim`, a process-oriented, discrete simulation library written in C++. Our goal is to provide a simulation library based on a widely used and efficient object-oriented programming language. Object-orientation has long been recognized to be very useful for modeling complex systems; it is

not a surprise that SIMULA was one of the first object-oriented programming languages. `libcppsim` implements a limited set of simulation primitives which can be learned quickly and can be immediately used for implementing simulation models. The provided general framework can be extended if necessary with more complex, high-level functionalities.

Efficiency was one of our main concerns. For that reason the pseudo-parallel simulation process system is not based on the threading model of the underlying Operating System (OS). Threads are not handled efficiently in every OS, as thread switching may incur in the same overhead as process switching. Moreover, complex simulation models made of tens or thousands of simulation processes would probably overflow the resources (process or thread table size) of many operating systems. We address this problem by implementing the *coroutine* primitive as a C++ class. Coroutines implements a simple, user-level cooperative multitasking environment; while this lacks the features of multitasking, it is exactly what is needed to implement process-oriented simulations.

This paper is organized as follows. Section 2 describes how coroutines and simulation processes are implemented in `libcppsim`, and the process scheduling facilities are illustrated in Section 3. Random variate generation is described in Section 4 and output data analysis in Section 5. In Section 6 we show a usage example of the library, and conclusions are reported in Section 7.

2 SIMULATION PROCESSES

Coroutines [Marlin, 1980] are a programming construct which can be useful for implementing process-oriented simulation languages and libraries. Coroutines are blocks of code with associated state. A coroutine can suspend itself by calling another coroutine; however, unlike traditional subroutine call in which the caller is suspended until the callee terminates, coroutines may be reactivated in any order.

SIMULA has coroutines as a built-in facility; unfortunately, the C and C++ programming languages, which are widely used and for which efficient compilers are available, do not provide coroutines natively.

In C/C++, each function is associated to a data structure called *activation record* which is stored in the runtime (LIFO) stack. The activation record contains informations about the state of the routine, such as the value of local variables and the return address used to resume the execution after a function call. When a function call occurs, a new activation record is created and put on the top of the stack. All variables locally defined by the called routine are stored on the newly created activation record. When a routine terminates, its activation record is pushed from the stack. LIFO handling of the activation records does not work anymore with coroutines, because the currently active coroutine may not be the one associated with the topmost activation record in the stack. This implies that the order of activation of coroutines is not given by the LIFO stack handling. Also, it is not possible to grow

the stack of all the routines whose activation records are not on at the topmost position.

This problem can be solved in different ways. The first approach is based on the “copy-stack implementation” described by Helsgaun [Helsgaun, 1999]. The stack of the currently operating coroutine is kept in the C++ runtime stack. When a coroutine suspends, the runtime stack is copied in a buffer associated with the coroutine. The buffer is allocated in the dynamic heap, so it does not interfere with the normal stack operation. A coroutine is resumed by copying the content of that buffer back to C++’s runtime stack and setting the program counter to the saved execution location.

A second approach consists of making use of the context handling facilities provided by most Unix SysV-like environments. The OS provides functions allowing user-level context switching between different threads of control within a process using the `getcontext()`, `setcontext()` and `makecontext()` functions. In this approach each coroutine has its own run-time stack in a pre-allocated block in the heap.

Both approaches have advantages and disadvantages. The “copy-stack” approach is the most portable, as it can be implemented in every modern OS which implements the `setjmp()` and `longjmp()` system calls (almost everyone does). However, it requires a copy of the C++ runtime stack to be saved and restored every time a coroutine passes control, which may cause a considerable overhead. On the other hand, the approach based on the context handling facilities is simpler to use and more efficient, as the stack content is not copied. This approach is, however, less portable as not every OS implements the required system calls. Also, the context-based approach does not allow the coroutine context to grow past the dimension defined when the context is allocated. The user is required to set the maximum dimension of the context in advance; choosing a too small buffer results in erroneous runtime behaviors caused by stack corruption which can be difficult to debug. A complete discussion is given in [Helsgaun, 1999], and detailed portability considerations of different context manipulation strategies are described in [Engelschall, 2000].

The `libcppsim` library implements both stack-handling mechanisms, and the user can choose which one to use at compile time. As a general rule, the “copy-stack” variant is useful when developing the simulation program, as stack corruption errors do not occur and it is possible to measure the maximum stack size required by the application; then, when the program is verified, production runs may be done with the more efficient context-based approach.

Once coroutines are available, it is easy to define simulation processes on top of them. Fig. 1 shows the class hierarchy related to simulation process implementation. A simulation process is represented by the `process` class, and is derived from the `coroutine` class; a simulation process is a coroutine whose state must be preserved across invocations.

Each simulation process has a unique identifier and a

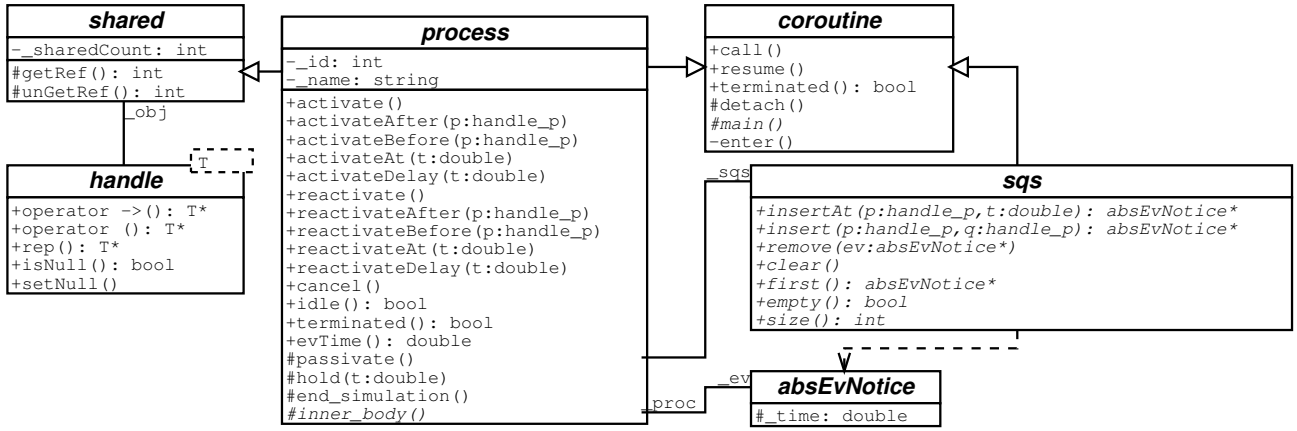


Figure 1: libcppsim Process Class Diagram

user-defined name. The process class has methods providing SIMULA-like scheduling primitives: it is possible to suspend the current process for a given amount of time, schedule the activation of another process in the future or cancel a pending process from the sequencing set. The actions performed by a simulation process are specified by defining the pure virtual method `inner_body()`. Fig. 2 shows the C++ interface of the process class.

The `handle<T>` class implements a “smart pointer” to an object of type `T`, where the class `T` must inherit from class `shared`. A smart pointer [Stroustrup, 1997] is a pointer to an object which keeps count of the number of times the object is referenced. When the last smart pointer is deallocated, the referenced object is automatically destroyed. Smart pointers are used for implementing a garbage collection mechanism. Our implementation of the `handle` class differs from the one proposed in [Stroustrup, 1997] as we store the reference count in the object rather than in the `handle`. This allows to preserve the counter even if the `handle` is converted to a C-style pointer and then back to an `handle`.

3 THE SEQUENCING SET

The Sequencing Set (SQS) is a data structure containing the list of simulation processes to be executed; the list is sorted by nondecreasing activation time order. Each simulation process contained in the SQS is associated with an event notice. The event notice basically contains the simulated timestamp at which the process is to be resumed, and references to that process. Additional informations may be present in order to make insertion or removal of event notices more efficient.

libcppsim provides two different implementations of the SQS data structure. The first is defined in class `sqsDll`, and is based on a doubly linked list of event notices. This data structure is very simple to implement, but insertions of event notices require linear time on average to be performed. A more efficient implementation is given by class `sqsPrio`, based on balanced search trees. The expected insertion time in this case is proportional to

the logarithm of the SQS size. More efficient data structures (e.g. calendar queues [Brown, 1988]) can be implemented by supplying a corresponding class, without other modifications being required.

4 RANDOM VARIATE GENERATION

Simulation programs require an efficient and statistically robust mechanism for producing streams of pseudo-random numbers with given distribution. libcppsim defines an abstract templated class `rng<T>` representing a pseudo-random number generator of numbers of type `T`. It is possible to generate streams of random integers, real or boolean values, by setting `T` to the appropriate datatype in a subclass.

It turns out that the basic ingredient for generating stream of pseudo-random numbers is a good uniform generator $RN(0, 1)$ over the interval $[0..1]$. We chose algorithm MRG32k3a from [L’Ecuyer, 1999], which is known to have long period and good statistical properties. We implemented different random variate generators using algorithms described in [Banks, 1998, Ch. 5]. These generators include the Uniform, Exponential, Gamma, Weibull, (Truncated) Normal, Bernoulli and Empirical distributions.

5 OUTPUT DATA ANALYSIS

Simulation results are sequences of observations for quantities of interest. Appropriate statistical techniques should be applied to analyze these sequences, as the observations are in general autocorrelated. libcppsim implements a set of classes dealing with collection of statistics; the relevant portion of the class diagram is shown in Fig. 3.

All types of statistical variables inherit from the `var<Tin, Tout>` abstract base class. This class represents “variables” which can be updated many times by feeding values of type `Tin`, and compute some function from the input data producing a result of type `Tout`. Each variable has two pure virtual methods: `update()` and `reset()`.

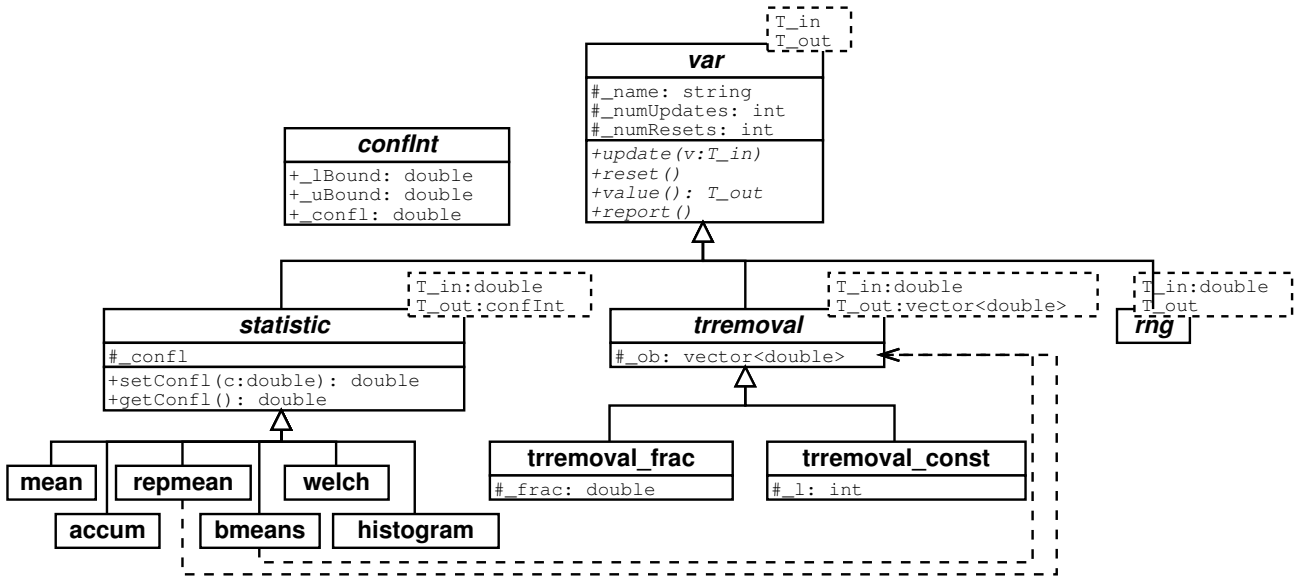


Figure 3: libcppsim Statistics Class Diagram

These are used to insert a new value and reset the variable to a known state. The result can be computed by invoking the `value()` method. Variables are not stateless; thus, it is possible that successive repeated invocations of the `value()` method return different results. This is useful for representing a random stream as a variable which is initially updated with the seed of the pseudo-random number generator, and whose `value()` method returns the next pseudo-random number from the stream, at the same time updating the seeds. Class `var` defines also a pure abstract `report()` method which can be used to display a report about the content of the class.

The statistics class represents a base class for a number of predefined statistic functions. In addition to the methods and attributes defined in its parent class `var`, statistics contains an attribute representing the confidence level of the computed result. Statistics are special kind of variables returning objects of type `confInt` (confidence intervals). The following statistics can be computed:

mean This class is used to compute a confidence interval for the mean of a sequence of *statistically uncorrelated* observations.

accum This class is used to compute the time-weighted sum of observations Y_1, Y_2, \dots, Y_n , $n > 1$ with timestamps respectively t_1, t_2, \dots, t_n . The result is $A = \sum_{i=1}^{n-1} Y_i(t_{i+1} - t_i)$

repmean Computes the mean of a sequence of observations using the method of independent replications. The simulation must be repeated R times, each run using different random streams and independently chosen initial conditions. Averages across replications can be used to obtain an accurate estimator of the sample mean [Banks, 1998].

bmeans The method of *batch means* divides the output data from one replication into a number of batches,

and then treating the means of these batches as if they were independent to compute the sample mean [Banks et al., 2001].

welch This class implements the graphical procedure of Welch [Welch, 1983] for identifying the length of the initial warm-up period. Values from R independent replications are collected. The values across different replications are first averaged and then smoothed by computing moving averages. The plot of the moving average becomes approximately constant after the warm-up period is over.

histogram This class computes an histogram profile from a sequence of values. The user supplies an expected lower bound and upper bound for the observed values, and the number of cells (bins) in which the histogram will be divided.

The `trremoval` class is used to model algorithms used for removing the initialization bias from a sequence of observations. The simulation run produces the sequence of observations Y_1, Y_2, \dots, Y_n for some parameter of interest. Using the whole sequence to compute the statistics is often dangerous, as there is a bias associated to artificial or arbitrary initial conditions. One method to overcome this limitation is to divide the sequence of observations into two phases: an initialization phase which includes observations Y_1, Y_2, \dots, Y_d , followed by a data-collection phase $Y_{d+1}, Y_{d+2}, \dots, Y_n$. Observations from the initialization phase are discarded, and only those in the data-collection phase are used for computing the required statistics. The parameter d , $0 < d \ll n$ is called *truncation point*.

Classes inheriting from `trremoval` are used to estimate the truncation point d . Currently, there is no automatic, general and correct method for detecting the length of the

```

typedef handle<process> handle_p;

class process :
  public coroutine, public shared {
public:
  process( const string& name );
  virtual ~process( );
  //// Modifiers ////
  // put in front of the sqs
  void activate( void );
  // activate AFTER q
  void activateAfter( handle_p& q );
  // activate BEFORE q
  void activateBefore( handle_p& q );
  // activate at time t
  void activateAt( double t );
  // activate with delay dt
  void activateDelay( double dt = 0.0 );
  void reactivate( void );
  void reactivateAfter( handle_p& q );
  void reactivateBefore( handle_p& q );
  void reactivateAt( double t );
  void reactivateDelay( double dt = 0.0 );
  // remove from the sqs
  void cancel( void );
  //// Accessors (const) ////
  // Is this process idle?
  bool idle( void ) const;
  // Is this process terminated?
  bool terminated( void ) const;
  // The next reactivation time
  double evTime( void ) const;
protected:
  //// Modifiers ////
  void hold( double dt );
  void passivate( void );
  void stop_simulation( void );
  //// Accessors ////
  // Currently executing process
  handle_p& current( void ) const;
  // Current simulation time
  double time( void ) const;
  // Simulation process' body
  virtual void inner_body( void ) = 0;
};

```

Figure 2: Interface of the process Class.

initialization bias. Different approaches have been proposed in the literature [White Jr. et al., 2000], although their effectiveness has not been agreed upon. For that reason, simulation packages usually implement the strategy of removing a prefix of fixed length (eg, 20%) from the sequence of observations; this approach is implemented by class `trremoval_frac`. This class can be used if the exact length of the warm-up period is known by the modeler.

6 EXAMPLE

We show in Fig. 4 a simple usage example of the `libcpcsim` library. The `main()` function performs a set of initialization tasks. The first is the creation of a new simulation context with an associated SQS object which is used by all processes. As in the original SIMULA language, it is possible to nest simulation contexts, allowing independent simulations to be executed inside other sim-

ulations. Once a simulation context has been created, it is possible to instantiate and schedule for activation one or more simulation processes.

```

#include <iostream>
#include "cpcsim.hh"

class job : public process {
public:
  job( const string& name ) :
    process( name ) {
    cout << "Created job '"
         << name << "' " << id() << endl;
  };
  virtual ~job( ) {
    cout << "Destroyed job '"
         << name() << "' " << id() << endl;
  };
protected:
  void inner_body( void ) {
    hold( 10 );
  };
};

class source : public process {
public:
  source( const string& name ) :
    process( name ) { };
  virtual ~source( ) { };
protected:
  void inner_body( void ) {
    handle<job> j;
    for ( int i=0; i<10; i++ ) {
      hold( 5 );
      j = new job( "job" );
      j->activateAfter( current() );
    }
    stop_simulation( );
  };
};

int main( void )
{
  simulation* sim = simulation::instance();
  // Instantiates a simulation context
  sim->begin_simulation( new sqsDll() );
  // Creates a simulation process
  handle<source> src = new source( "source" );
  // Schedules the process for execution
  src->activate();
  // Runs the simulation
  sim->run();
  // Cleans up the SQS
  sim->end_simulation();
  return 0;
};

```

Figure 4: Usage Example of the `libcpcsim` Library

7 CONCLUSIONS

In this paper we described a SIMULA-like process-oriented simulation library implemented in C++. The library is based on the coroutine abstraction implemented as a C++ class, and has been successfully compiled on various versions of the Linux OS and other flavors of Unix including Digital/Compaq OSF and NetBSD. The library provides facilities such as simulation processes, random variate generators and basic statistical functions. The `libcpcsim` library has been used to implement UML- Ψ , a process-oriented simulation tool for performance evaluation of UML specifications [Balsamo and

Marzolla, 2003; Marzolla, 2004].

Due to its object-oriented nature, the library can be easily extended with additional functionalities. We are currently improving the library in several directions. First, we are implementing more high-level synchronization mechanisms (such as semaphores and mailboxes) which can be useful for developing certain complex simulation models. We are also implementing more statistical functions in order to provide the modeler with a set of more advanced mechanisms for evaluation of the simulation results.

Availability

Full source code of the `libcpps` simulation library is available on the author's web page at <http://www.dsi.unive.it/~marzolla>

Acknowledgments

This work has been partially supported by MIUR research project FIRB "Performance Evaluation of Complex Systems: Techniques, Methodologies and Tools".

REFERENCES

- Balsamo, S. and M. Marzolla (2003, June 9–11). Simulation modeling of UML software architectures. In D. Al-Dabass (Ed.), *Proc. ESM'03, the 17th European Simulation Multiconference*, Nottingham, UK, pp. 562–567. SCS–European Publishing House.
- Banks, J. (Ed.) (1998). *Handbook of Simulation*. Wiley–Interscience.
- Banks, J., J. S. Carson, B. L. Nelson, and D. M. Nicol (2001). *Discrete-Event System Simulation*. Prentice Hall.
- Birtwistle, G. (1979). *DEMOS—A system for discrete event modelling on Simula*. MacMillan Press.
- Brown, R. (1988, October). Calendar queues: A fast 0(1) priority queue implementation for the simulation event set problem. *Comm. of the ACM* 31(10), 1220–1227.
- Dahl, O.-J. and K. Nygaard (1966, September). SIMULA—an ALGOL-based simulation language. *Comm. of the ACM* 9(9), 671–678.
- Darmont, J. (2000, January). DESP-C++: a discrete-event simulation package for C++. *Software–Practice and Experience* 30(1), 37–60.
- Engelschall, R. S. (2000, June 18–23). Portable multithreading—the signal stack trick for user-space thread creation. In *Proc. USENIX Ann. Tech. Conference*, San Diego, CA.
- Helsgaun, K. (1999). A portable C++ library for coroutine sequencing. *Datalogiske Skrifter (Writings on Computer Science)* 87, Roskilde University, Denmark.
- Hlavička, J., S. Racek, and P. Herout (1999, December). C-Sim v.4.1. Research Report DC-99-09, DCSE CTU Prague Publishing, Prague.
- L'Ecuyer, P. (1999, January). Good parameters and implementations for combined multiple recursive random number generators. *Operations Research* 47, 159–164.
- Little, M. C. JavaSim Home page. <http://javasim.ncl.ac.uk/>.
- Little, M. C. and D. L. McCue (1993, July). Construction and use of a simulation package in C++. Computing Science Tech. Report 437, University of Newcastle upon Tyne.
- Marlin, C. D. (1980). *Coroutines: A Programming Methodology, a Language Design and an Implementation*, Volume 95 of LNCS. Springer-Verlag.
- Marzolla, M. (2004, February). *Simulation-Based Performance Modeling of UML Software Architectures*. PhD thesis, TD-2004-1, Dipartimento di Informatica, Università Ca' Foscari di Venezia, Italy.
- Schwetman, H. (2001). CSIM19: a powerful tool for building system models. In *Proc. 33rd Winter Simulation Conference*, pp. 250–255. IEEE Computer Society.
- Stroustrup, B. (1997). *The C++ Programming Language* (3rd ed.). Addison–Wesley.
- Varga, A. (2001, June 6–9). The OMNET++ discrete event simulation system. In *Proc. ESM'01, the 15th European Simulation Multiconference*, Prague, Czech Republic, pp. 319–324. SCS–European Publishing House.
- Welch, P. D. (1983). The statistical analysis of simulation results. In S. Lavenberg (Ed.), *The Computer Performance Modeling Handbook*, pp. 268–328. Academic Press.
- White Jr., K. P., M. J. Cobb, and S. C. Spratt (2000). A comparison of five steady-state truncation heuristics for simulation. In *Proc. 32nd Winter Simulation Conference*, Orlando, FL, pp. 755–760. SCS International.

AUTHOR BIOGRAPHY

MORENO MARZOLLA received MSc and PhD degrees in Computer Science in 1998 and 2004 respectively, both from "Ca' Foscari" University of Venice. From 1999 to 2001 he was a software engineer at the Italian National Institute for Nuclear Physics. He currently holds a post-doc position at the University of Venice. His research interests include performance evaluation of software systems, distributed systems monitoring and parallel and distributed simulation.