

Simulating Process Chain Models with OMNeT++

Falko Bause, Peter Buchholz, Jan Kriege, Sebastian Vastag
Technische Universität Dortmund, Informatik IV
D-44221 Dortmund, Germany

{falko.bause,peter.buchholz,jan.kriege,sebastian.vastag}@udo.edu *

ABSTRACT

This paper presents an approach to simulate complex hierarchical process chains resulting from large logistics networks in *OMNeT++*, a discrete event simulation environment designed for communication networks. For this purpose *OMNeT++* has been integrated as a new simulation engine into the *ProC/B* toolset which is designed for the analysis and optimization of large logistics networks. The paper highlights the main steps of the automatic transformation of a hierarchical process chain model into a hierarchical model in *OMNeT++*. Furthermore it shows how the transformation has been validated and how detailed performance figures can be evaluated with *OMNeT++*.

Categories and Subject Descriptors

I.6.5 [Simulation and Modeling]: [model development]

General Terms

Design, Performance

Keywords

OMNeT++, DES, *ProC/B*, Logistics

1. INTRODUCTION

For the development and operation of contemporary networks in logistics, model based analysis and in particular the use of discrete event simulation is becoming an important factor to ensure that the networks meet the requirements concerning technical measures like delivery times or service levels and, on the other hand, are also cost effective. In the past the different workflows in a logistics network have been specified with process chains as a poorly descriptive tool that does not allow one to derive simulation models from

*This research was supported by the Deutsche Forschungsgemeinschaft as part of the Collaborative Research Center "Modelling of Large Logistics Networks"(559).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIMUTools 2008 March 3-7, 2008, Marseille, France.
Copyright 2008 ACM ISBN 978-963-9799-20-2 ...\$5.00.

the description. This, however, implies that required simulation models have to be specified on their own without any formal relation to the process chain model. Of course, this approach has the disadvantage that different models have to be created for one system with all the known problems of additional modeling effort or inconsistencies between the models. Thus, the use of the entire process chain model as a base model for a detailed simulation model of a logistics network is highly recommendable.

To realize this approach partially informal process chain models have to be enhanced by formal information necessary to build a simulation model and adequate software tools for simulation have to be available. Of course, simulation of process chains is not a new idea [18], but a general approach which allows one to refine a high level process chain into a detailed simulation model and which can cope with the complexity and size of models of today's logistics networks is still missing. Available simulation tools for this purpose are either restricted prototypes [17] or extensions of business process modeling tools [11]. In both cases the capabilities of representing and analyzing more complex models is limited. Available simulation tools for manufacturing systems [15, 16] that have been developed for large systems lack basic features which are necessary to model logistics networks and general simulation frameworks are too low level such that an adequate modeling of complex process chain models requires too much effort.

In the past we developed a class of hierarchical process chain models which include the necessary information to map them onto discrete event simulation models [1]. The model class, denoted as *ProC/B*, is based on a hierarchical description where activities of a process chain are performed by some function unit which itself can be a complex process chain or some basic unit describing the consumption of space or time. The resulting models may include an arbitrary number of hierarchical levels which form an acyclic graph. Originally, *ProC/B* models have been mapped onto simulation models using the tool HIT [3] which has been developed in the mid eighties for modeling complex computer and communication systems. HIT perfectly supports the hierarchical structure of *ProC/B* models and allows the analysis of results according to arbitrary paths through the hierarchy which is an important feature, in particular, if economic measures where cost drivers become important should be evaluated via simulation. However, the use of HIT also introduces some serious limitations. Since HIT generates a simulation model in the language SIMULA, a runtime environment for SIMULA has to be available to run a sim-

ulation. Unfortunately, the number of available SIMULA compilers is rather limited. Furthermore, HIT as a nearly 20 year old tool does not support several modern features of an object oriented simulation environment like animation or interfaces to software tools for post-processing of results or for the administration of models. For these reasons we decided to integrate a new simulation tool in our modeling environment and to support a mapping of *ProC/B* models onto the corresponding models.

An adequate simulation platform has to observe the following requirements:

- The full *ProC/B* model world has to be mapped in the simulation model.
- The hierarchy of the process chain models has to be adequately represented in the simulation model.
- Detailed measures that are definable in *ProC/B* should be analyzable in the resulting simulation models.
- The simulation tool has to be driven by the *ProC/B* interface.
- Simulation should be easily made interoperable with other tools of the *ProC/B* environment like the optimization tool OPEDo [5] or the trace analyzer Traviando [13].
- The simulation environment should be stable, should allow the definition and simulation of large models and it should support modern features of object oriented simulation.
- The simulation environment should be freely available for research according to some adequate open source license.

The last two points restricted the number of available tools significantly since most available open source simulation tools were not adequate to really simulate, in an efficient and error-free way, large models as they result from large logistics networks. After a more detailed look on the remaining tools, our choice was *OMNeT++* [14], a simulation environment generating simulations in C++. Although *OMNeT++* had been developed and used for communication systems, it is well suited for the mapping of hierarchical process chain models and it also fulfills almost all of the above requirements. Nevertheless, the mapping of *ProC/B* models onto *OMNeT++* is far from being trivial since complex hierarchies have to be transferred from one view into the other.

This paper introduces the combination of *ProC/B* and *OMNeT++* to build a new and powerful simulation environment for process models of logistics networks. The work presented is partly based on [10, 20]. In the following section we briefly present the basic tools *ProC/B* and *OMNeT++*. Afterwards, in section 3 it is shown how the hierarchical structure of *ProC/B* is mapped onto a module structure of *OMNeT++*. Then we show how the behavior of *ProC/B* processes is performed in *OMNeT++*. Section 5 is devoted to the validation of the mapping, followed by a first comparison of *OMNeT++* and *HIT* by means of small examples. The paper ends with the conclusions.

2. BASIC SOFTWARE TOOLS

The approach we present in this paper uses *ProC/B* as input format, maps the models to *OMNeT++* and simulates the resulting model using *OMNeT++*. In this section we briefly present the main features of the *ProC/B* approach and give afterwards a very brief introduction into *OMNeT++*. For further details about both tools we refer to the literature [1, 14].

2.1 Introduction to ProC/B

ProC/B [1] is a process chain-based modeling approach which is used in the collaborative research center “Modelling of Large Logistics Networks” 559 (CRC 559;[7]) for modeling and performance evaluation of logistics networks. *ProC/B* accounts for the specifics of the application area by capturing the structure in form of function units (FUs) and the behavior by process chains (PCs). In *ProC/B*, FUs might offer services, which can be used by activities of process chains. Each service is again described by a process chain.

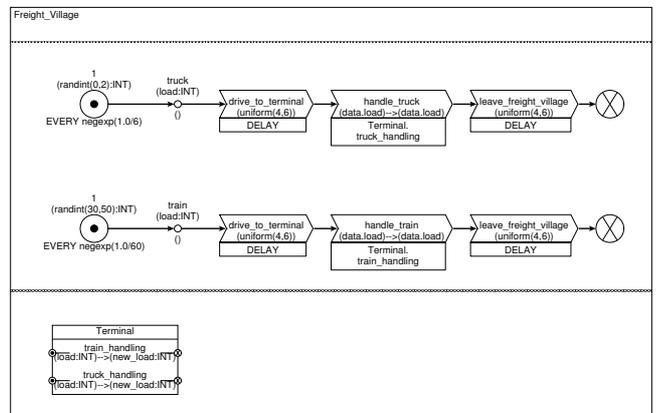


Figure 1: *ProC/B* top level model “Freight Village”

Figs. 1 and 2 present an example of a *ProC/B* model representing a simplified freight village. A freight village is a node of a logistics network which provides facilities for storing goods temporarily and transshipment between several type of carriers.

The top level of the model (see Fig. 1) is specified by FU `Freight_Village` whose behavioral part is described by two PCs: `truck` and `train`. The structure part consists of a single (user defined) FU, named `Terminal`, which offers two services: `truck_handling` and `train_handling`. Services can be compared to functions in programming languages. In the example both services have an input parameter (`load`) and an output parameter (`new_load`). Behavior and structure part of a FU specification are interrelated by expressing which service of which FU performs an activity. In Fig. 1 the two PCs `truck` and `train` consist of three process chain elements (PCEs) each, and in both cases the second activity calls a service of FU `Terminal`. The inner view of FU `Terminal` is shown in Fig. 2. The offered services are specified by PCs and some of their activities use the services of two so-called standard function units which offer predefined services (e.g. `request` and `change`). *ProC/B* offers two kinds of standard FUs: servers and storages. Servers (see `fork-`

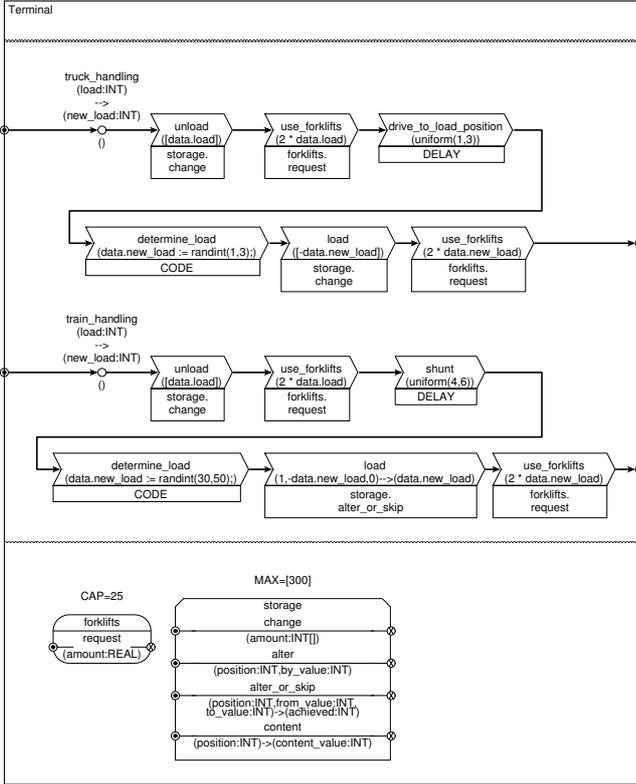


Figure 2: Function Unit “Terminal”

lifts in Fig. 2) capture the familiar behavior of traditional queues describing the consumption of time and storages describe the consumption of space (see `storage` in Fig. 2) and support the manipulation of passive resources. A simplified version of a storage is a so called counter, which is a standard FU often used for modelling synchronization aspects. A change to a counter or a storage is immediately granted iff the result respects specified upper and lower bound vectors; otherwise the requesting process gets blocked until the change becomes possible.

Process chains directly visualize behavior. The freight village model of Figs. 1 and 2 reads as follows: Incarnations of process chain `train` are generated according to a Poisson distribution (with a mean of 60 time units). Each train has a load which is initially chosen by random according to an uniform distribution (between 30 and 50). After incarnation, the train “drives” to the terminal which is modeled here by a delay of the process for a uniformly distributed duration. Afterwards the train “is handled” by service `train_handling` of `Terminal`. This might result in a change of the train’s load. Finally the train “leaves” the freight village and the process terminates at the sink. Considering Fig. 2 we see that handling a train means first to unload the train, which is possible if the `storage`’s capacity of 300 units is not exceeded, otherwise the train has to wait until unloading is possible. Afterwards the server `forklifts` is called, which is a multi-server queue with 25 servers and a (default) FIFO service strategy. The service time for the requesting process is determined by the expression `2 * data.load` thus model-

ing the time for unloading a loaded train¹. Afterwards the train “shunts” to a new position (which is again modeled by a delay of the process) and determines the new load. The new load is removed from the storage if possible, otherwise the available number of units are removed from the storage (which is the semantics of service `alter_or_skip`). Finally service `request` of the server `forklifts` is called again before the process “leaves” the terminal. The behavior of PC `truck` and service `truck_handling` reads similarly. In the sequel we will use the term process for the process description and its incarnations.

In the course of the CRC 559 a toolset has been developed which provides a graphical user interface to specify *ProC/B* models and transformer modules which map *ProC/B* models to the input languages of existing analysis tools, so that *ProC/B* models can be analyzed automatically (cf. [1] and Fig. 3).

Simulation is often applied for a detailed analysis, since it is applicable to all *ProC/B* models. In the past, simulation was only supported by *HIT* [3, 4]. *HIT* is a modeling environment which does not only provide a simulator, but also offers efficient non-simulative analysis algorithms being based on product-form queueing networks. *HIT* is basically tailored to steady-state analysis, based on a single replication approach. Time Series Analysis techniques are applied to individual streams of data produced by the simulation. A key feature of *HIT* is that these streams may be itemized in detailed ways. E.g., in Fig. 2 it might be of interest to measure separately the number of service calls of FU `storage` caused by trucks and trains. *HIT* provides facilities to describe and to evaluate measures for such activities at a lower level which are caused by some higher level originator and to itemize corresponding results with respect to the originators. As mentioned, *HIT* is nearly 20 years old and needs a *SIMULA* compiler for execution. Therefore we recently integrated *OMNeT++* into the *ProC/B* toolset trying to benefit from the features of a modern object-oriented simulation environment.

2.2 OMNeT++

OMNeT++ is a public-source simulation environment that has been developed for the modeling of communication protocols and has been extensively used in this area. Although it is mentioned on the web page [14] that *OMNeT++* has also been used for the analysis of business processes there is nothing available about this application and it does not seem that a complete mapping of hierarchical business processes onto *OMNeT++* models has been done before. The whole tool environment includes a graphical front end and several other tools that support the modeling and simulative analysis of complex systems. Of particular interest is the simulation kernel which is written in C++ and offers several classes to support the specification of complex hierarchical models. Furthermore, the resulting simulation models are known to be rather efficient.

The basic entities of an *OMNeT++* simulation are modules. Modules can be simple, which means that they are implemented as C++ classes, or compound modules which means that they are composed of other simple or compound

¹Access notations to parameters and variables of processes are prefixed with keyword `data` for technical reasons in order to distinguish them from global variables. Global variables are not shown in Figs. 1 and 2.

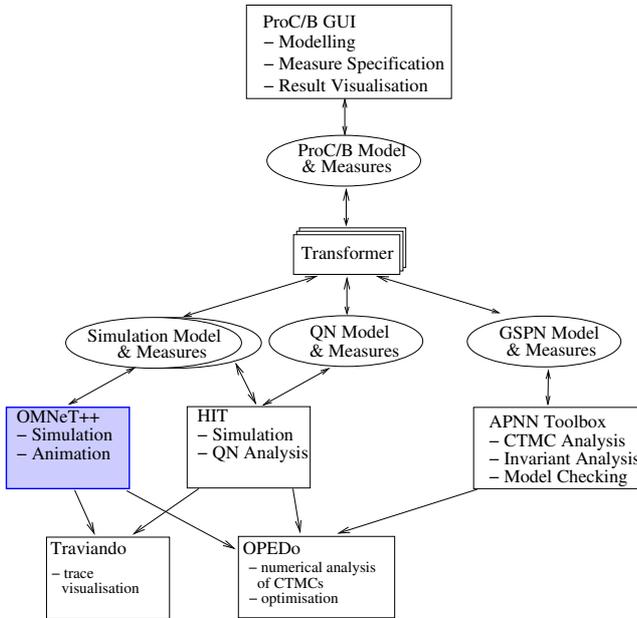


Figure 3: *ProC/B* toolset

modules. In this way *OMNeT++* models are hierarchical. The complete model containing the overall hierarchy is denoted as the system module. Modules communicate via gates using messages. Gates can be input or output and a module may have an arbitrary number of gates. Messages are sent either directly to a gate or along a path. Basically paths are used to describe the transfer of messages over some medium. Therefore they offer parameters to specify e.g. the bandwidth or loss rate. The connection of modules via paths is specified in the **.ned* file which includes the structure of the model and can be defined with the help of a graphical interface. The graphical interface can also be used for the animation of the running model by visualizing messages that are sent along a path from one module to another.

An arriving message is interpreted in a module as an event and the user has to specify a routine *handleMessage()* for each arriving message type. Messages themselves can be structured data types and may include information that is used in the corresponding *handleMessage()* routine. In the routines new messages may be generated immediately or after some delay and already scheduled messages may be deleted. Thus, the basic event driven approach is realized by the processing and sending of messages. Apart from this general mechanism, the simulation kernel of *OMNeT++* offers a lot of support to realize complex simulation models like up to date random number generators, support for statistical evaluation of results or support for parallel replications.

From this very brief description it should become clear that both, *ProC/B* and *OMNeT++* use a hierarchical structure to describe models. However, at a second view it becomes clear that the model views differ in several yet important details. *OMNeT++* has been designed with communication systems in mind such that messages have a physical meaning whereas in *ProC/B* hierarchy is introduced by calling services of FUs without explicit messaging. Another important aspect is, as already mentioned, the definition of

detailed and origin dependent measures which are not directly supported by *OMNeT++* and therefore have to be implemented separately.

The challenge is to get a correct mapping from *ProC/B* onto *OMNeT++*. Since correctness of the mapping cannot be formally defined because only a subset of *ProC/B* has a formal semantics in form of a Petri-net mapping [6], we define correctness by comparing the simulation using *HIT* and *OMNeT++*. The *HIT* simulation is usually taken as the correct behavior, since we defined an operational semantics of the whole *ProC/B* paradigm via *HIT* [2], and thus the *OMNeT++* model has to show the same behavior. Of course, a detailed comparison implies that the model is completely deterministic since otherwise different random number streams will necessarily result in different behaviors such that only statistical results can be compared using adequate statistical methods [12]. We used both, simple deterministic models to show that the basic behavior is the same and more complex stochastic models to compare statistically the result measures.

In the following two sections we first describe how the hierarchical structure of a *ProC/B* model is mapped onto a corresponding structure of an *OMNeT++* model. In the subsequent section the mapping of the behavior is presented. Both steps are accompanied by small examples showing the basic ideas.

3. MAPPING OF STRUCTURE

Our implementation of mapping *ProC/B* to *OMNeT++* consists of two main components: The converter *procb2ned* and a library named *Osimu* containing implementations of *ProC/B*'s behavior as modules for *OMNeT++*.

ProC/B models are specified graphically in the *ProC/B* editor and are stored in files. Next to the model itself the editor allows for saving experiment descriptions in separate files. As stated above, these general model/experiment description files can be used in different analyzers, either numeric or simulative. Thus, the generic model descriptions generated by the editor have to be translated to specific input formats. For our simulative analyzer based on *OMNeT++* the converter *procb2ned* reads process chain models and outputs *OMNeT++* network descriptions (**.ned*-files) as a direct input format for the *OMNeT++* simulation system. As *OMNeT++* supports hierarchical modeling of modules, *ProC/B*'s structure is preserved and included in **.ned*-files by *procb2ned*.

OMNeT++ requires behavior to be located in atomic entities called "simple modules". They are written in C++ and handle arriving messages in order to trigger specific reactions. The sources are combined with **.ned*-files with an identical naming scheme, describing module interfaces to *OMNeT++*'s simulation system. For model design these basic modules are instantiated and related by connections in another **.ned* file, forming a "compound module". Non-basic models can be used similar to basic modules, making it easy to form hierarchical models.

As *ProC/B* also supports hierarchical modeling, *procb2ned*'s primary task is to map given hierarchies. For every language element in *ProC/B* exists a corresponding implementation as a basic module in *OMNeT++*. PCEs only form linear structures at the same model level, so *procb2ned* simply inserts them as basic modules into its output **.ned* files. Of course, synchronization and event driven generation of new

processes are also possible in *ProC/B* using available language elements (see [1]) which are realized as C++-implementations in *OMNeT++*. Following the rule of one module per *ProC/B* language element, Standard-FUs like ServerFU, StorageFU and CounterFU are also inserted directly into the model.

Hierarchies in *ProC/B* are the domain of constructed FUs, usually including at least one process chain offered as a service by the FU. To model these FUs in *OMNeT++*, compound modules are formed by **.ned* files, one for each constructed function unit. If the converter reads a constructed FU on input, it goes one level down in the recursion, applying the above mapping rules to a new **.ned* file named after its FU in *ProC/B*. After returning from recursion, the compound module representing a constructed function unit can be used like any other basic module.

The subsequent step is to map process flow through a process chain by establishing connections between modules. The acting entities of *ProC/B* are all processes within a module. While processes are no specific object-types in HIT, it was a natural choice to map exactly one process type to exactly one message type in *OMNeT++*. Hence, *ProC/B*'s connections between process chain elements are mapped to module connections in *OMNeT++*.

As shown in Fig. 1, connections in *ProC/B* only exist within a PC specification, there are no explicit connections from PCEs to FUs or the other way round. Only implicit relations between PCEs and FUs exist by specifying parameters in PCEs denoting which FU and offered service they call.

We transferred this idea to *OMNeT++* by using traditional message passing through gateways and direct message sending as two separate forms of connections. The first task is done straightforwardly by *procb2ned*. Basic modules act as PCEs (excluding sinks) and obtain one connection to their successor, forming a structure similar to process chains in *ProC/B*. For this purpose, every PCE module has at least one set of input/output gateways acting as a socket for *OMNeT++*'s connections.

Again, relationships between modules of PCEs and FUs exist only implicitly in *OMNeT++*. Two parameters are given in the **.ned* file for every instance of a PCE using a function unit: The identifier of the FU and the name of the offered service (keeping in mind that FUs can offer multiple services). This information is used in the PCE initialization phase to find the reference to their loosely bound function unit. Requesting a service in the simulation phase is done by transferring messages directly using *OMNeT++*'s `sendDirect()` method to the function unit bound to the service. Finishing a FU's service is also signaled by returning the message.

Using *OMNeT++*'s alternative way to transfer messages has some advantages compared to the traditional way of using module connections:

1. As FUs can be used by possibly infinitely many PCEs, omitting explicit connections helps to keep models concise. The target to which direct messages are sent to is determined during the initialization phase by *OMNeT++* and saved as a reference, so no extra time is consumed when analyzing the model.
2. Messages sent directly keep track of their senders on a stack, so returning a process message to its sender

after performing a service is a simple task in the FU's implementation.

3. The visual appearance in the *ProC/B* editor and *OMNeT++/TKenv* is kept similar.

Function units need one input gate per offered service. Output gates are redundant here, as the virtual sink terminating the service's process chain will return the message via direct transfer to the calling PCE. The ability to send and receive direct messages requires some preconditions for modules in *OMNeT++*: Direct messages can only be delivered to dedicated input gates without any other incoming connection. Therefore, PCE modules calling FUs need an additional input gate reserved for callbacks of their associated FUs. By convention, new processes arrive at the first, status messages from FUs at the second input gateway.

3.1 Animation

An important additional benefit of using *OMNeT++* for simulating *ProC/B* models is the animation capability of *OMNeT++*'s graphical workbench *OMNeT++/TKenv*. Existing *ProC/B* analyzers are tuned according to efficiency and performance of the solution and are consequently batch processing systems, making it difficult to explain the dynamic behavior of processes from the model. However, such an explanation is often important in teaching and also in real projects as we noticed when modeling large systems in cooperation with real users. Using *OMNeT++/TKenv*, messages moving between modules can be animated by a moving red dot as an adequate visualization for processes moving through process chains and making use of FUs. In this way, the dynamics of a system is clearly visible.

ProC/B's graphical representation was carried to *OMNeT++*, using *OMNeT++*'s feature to define pictograms for modules. *procb2ned* assigns bitmaps to every instance of the basic module matching its type in *ProC/B*.

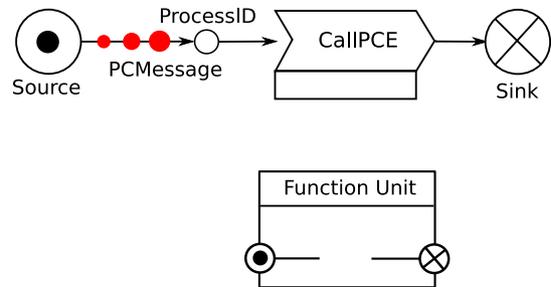


Figure 4: Animation of PCMessages

When setting *OMNeT++* to a "slow" running mode in *OMNeT++/TKenv*, the user can trace processes created by a source as a red dot moving along the process chain element's outgoing connection (see Fig. 4). Arriving at a PCE, the dot is delayed until the PCE ends its call to an FU. Since no permanent connection exists between PCEs and FUs, a temporary connection is drawn acting as a path for messages performing a request by being sent directly to the module of the FU (Fig. 5). When the request is served, a message can be seen moving backward to the calling PCE on a reverse connection. Furthermore, different windows can be opened

to view the animation simultaneously at different levels, i.e. in different FUs.

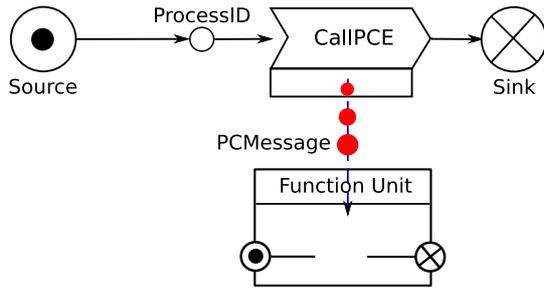


Figure 5: PCMessage moving over temporary connection

4. MAPPING OF BEHAVIOR

The second component of our framework is the mapping of the behavior of *ProC/B* language elements onto modules in *OMNeT++*. These modules are completely configurable by passing parameters, so we were able to compile a library *Osimu* containing generic implementations for PCEs and FUs. Every element of *ProC/B*'s language is mapped onto exactly one module implementation, inheriting into *OMNeT++*'s modeling interface by subclassing `cSimpleModule` directly.

OMNeT++ offers two different programming styles, lightweight processes and a transaction based programming paradigm. We decided to follow the latter as it matches the basic ideas of process chain models and it scales much better for large models. Since process chains are characterized by the interval of time consumed by a process between entering and leaving a chain, a transaction based discrete event approach makes it natural to map these intervals to arrivals and departures of messages to/from modules in *OMNeT++*.

The dynamic part of *ProC/B* models are processes following the route defined by process chains. Analogously, processes are represented by messages in *OMNeT++*. The mapping to *OMNeT++* is done by subclassing `cMessage` only once to `PCMessage` (short for "process chain message"). In *ProC/B*, transitions of processes between PCEs are instantaneous, time is consumed by requesting services at function units or by dedicated delay PCEs. This idea is reproduced in *OMNeT++* by messages of type `PCMessage` that use connections in zero time, leaving progress of model time to the modules.

Now we will describe two examples to explain our mapping of behavior more explicitly.

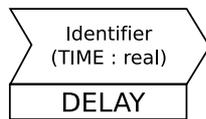


Figure 6: *ProC/B*'s symbol for a delaying PCE

Figure 6 shows the symbol of a Delay PCE as a simple language element of *ProC/B*. It has the task to delay arriving processes by some amount of time, either deterministic or by

Listing 1: Source of DelayPCE.ned

```
simple DelayPCE
parameters:
  delay : string
gates:
  in: in;
  out: out;
endsimple
```

Listing 2: Source of DelayPCE.cc

```
#include "DelayPCE.h"

void DelayPCE::handleMessage(cMessage* msg) {
  double delay = (double) param("delay");
  sendDelayed(msg, delay, "out");
}
```

a random number from some predefined distribution. Listing 1 is the according **.ned*-file used by *OMNeT++* to pass parameters to modules and define gateways where messages arrive and leave.

A process chain element DelayPCE is defined with just one pair of gates because it is connected to only two other process chain elements. Parameter `delay` is set by the **.ned*-file instantiating this module as specified in the original *ProC/B* model. Implementations of DelayPCE subclassing `cSimpleModule` have to respect the fact that a second or third or an arbitrary number of processes can arrive while the first process is still delayed.

The concrete implementation in listing 2 is short, overloading the function `handleMessage(cMessage* msg)` and implementing the specific reaction on an arrival of messages.

The delay time is generated from the predefined distribution, the following line in the code delays the incoming message by the amount of time using *OMNeT++*'s `sendDelayed()` method. Please note that no messages are stored inside the module, allowing the module to accept an infinite number of processes. Method `param()` reassembles the well known method `par()` to fetch parameters stored in **.ned* files. The new method parses expressions used in *ProC/B* models, which can be either arithmetic expressions or a *ProC/B* specific naming of random distributions.

SPEED : number
DIS : string
CAP : number

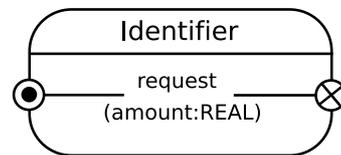


Figure 7: Server Function Unit in *ProC/B*

As an advanced example ServerFU is shown in Fig. 7. It represents a set of limited and identical resources which processes can request and use. ServerFU offers the service "request" to PCEs, a parameter for the requested amount of service has to be passed with the calling message.

Listing 3: ServerFU.ned

```
simple ServerFU
  parameters:
    inPath: bool,
    speed: numeric,
    capacity: numeric,
    discipline: string;
  gates:
    in: in;
    out: out;
endsimple
```

Listing 3 contains the definition of ServerFUs in *OMNeT++*'s modeling language. Three parameters are passed to the module by *OMNeT++* at runtime:

capacity number of resources offered by the server

speed of a resource to execute service calls. This means that calling PCEs request an amount of service according to some average resource. The concrete FU can be faster ($\text{speed} > 1$), slower ($\text{speed} < 1$) or an average resource ($\text{speed} = 1$).

discipline the resource scheduling the server uses.

By offering a single service, ServerFU only needs one pair of gates. As stated before, messages are delivered to these gates by direct send calls, making it unnecessary to connect this module with other elements in the model.

The parameter discipline plays an important role in *ProC/B* models, as the behavior of function units are matched to the way resources are shared in the real system. Currently three scheduling disciplines can be mapped from *ProC/B* to *OMNeT++*:

FCFS queues serve requests for resources by the rule "first come, first served". Processes which obtained a resource, allocate it according to the amount defined by the constant or distribution of the service call and parameter "speed".

IS "Infinite Server", every request is immediately granted and takes the time specified by the service call and parameter "speed".

PS "Processor Sharing", all requests are immediately granted. Every process makes use of the full set of resources (accelerating service time by the number of resources), but has to share resources with other process using the server at the same time. Capacity is distributed uniformly among all processes (slowing down execution time by the reciprocal value of the number of processes).

It is possible to interpret these three types of scheduling as different kinds of servers, yet their implementation in *OMNeT++* uses only one module to simplify the structure mapping by *procb2ned*. Internally ServerFU makes use of the strategy pattern [9] to vary its behavior according to parameter "discipline".

An excerpt of ServerFU's source is shown in listing 4. The `handleMessage()` method is divided into two parts by an `if` clause, newly arriving process messages with `selfMessage` set to false are served in the lower part.

Listing 4: Excerpt from ServerFU.cc

```
void ServerFU::handleMessage(cMessage* msg) {
    PCMessage* message = (PCMessage*) msg;

    if (message->isSelfMessage()) {
        // message was scheduled by handleProcess()
        discipline->handleSelfMessage(message);
        finishService(message);
        dismissMessage(message);
        return;
    }

    welcomeMessage(message);
    writeNameInPath(message);

    discipline->handleProcess(message);
}

void ServerFU::FCFSDiscipline::handleProcess(
    PCMessage* message) {
    if (parent->numberMsgInSystem <= parent->
        serverCapacity) {
        parent->scheduleAt(simulation.simTime() +
            serviceTime(message), message);
    }
    else fcfsQueue.insert(message);
}

void ServerFU::FCFSDiscipline::handleSelfMessage(
    PCMessage* message) {
    if (fcfsQueue.empty()) return;

    PCMessage* msgFromQueue = (PKMessage*) fcfsQueue.
        pop();
    parent->scheduleAt(simulation.simTime() +
        serviceTime(msgFromQueue), msgFromQueue);
}

double ServerFU::FCFSDiscipline::serviceTime(
    PKMessage* message) {
    return parent->getTaskTime(message) / parent->
        stdSpeed;
}

void ServerFU::ISDiscipline::handleProcess(PCMessage*
    message) {
    parent->scheduleAt(simulation.simTime() +
        serviceTime(message), message);
}
```

At the beginning, `welcomeMessage` computes some basic statistics of arriving processes as described in section 4 and increments `numberMsgInSystem`. Additionally, the ServerFU pushes its name in `writeNameInPath` on the stack keeping track of every process chain element the message passed through. The next line is part of the strategy pattern, `discipline` holds objects of type `ServerFU::Discipline` encapsulating FCFS, Infinite Server or Processor Sharing as described above. Those strategies are instantiated depending on parameter `discipline` in listing 3 on module's initialization. Their behavior on newly arrived processes is specified in `handleProcess()`. Here we present the methods of FCFS and IS as examples: In FCFS, time is granted to processes as long as the servers capacity is not exceeded. Otherwise the process is enqueued until more resources become available. The Infinite Server is even more simple, it just accepts every process.

In both examples, time consumption is modeled by scheduling process messages to the function unit itself, adding the amount of time the service will take before sending. When

the message returns, `selfMessage` is true and the upper part of `handleMessage()` is executed. Again, an object of type `ServerFU:Discipline` handles processes a second time. In FCFS, the first process message waiting for free resources is removed from the queue and immediately scheduled for completion of the service. The corresponding method for Infinite Server is not shown here, it equals to an empty implementation since this discipline does not need any further action after completing a service.

4.1 Translation of Result Measures

The main focus when analyzing a simulation model is on determining quantitative results for the model, like for example throughputs or response times. *ProC/B* offers the possibility to measure properties at every FU, though depending on the type of the FU the available properties may differ: Throughput, response time and population can be measured at any FU. Additionally, for every server the utilization and for every storage the state can be examined. For composed FUs the modeler may define further measures (called rewards in *ProC/B*). *ProC/B* allows for three different types of rewards: *Count*, *event* and *state*. Rewards of the type *event* can be used for serially collecting values, rewards of the type *count* for estimating rates and rewards of the type *state* for the description of trajectories. Those types are used for the realization of the standard measures like throughput or response time as well. While the user-defined measures have to be updated manually (*ProC/B* provides a model element for updating those rewards), the standard measures are updated automatically whenever a process enters or leaves a FU.

As already mentioned *ProC/B* allows for streams to be itemized in detailed ways. This enables for example the measurement of the train population at the terminal in Fig. 2 without counting trucks. To achieve this, the modeler can specify a path consisting of elements in the *ProC/B* model. Only processes, that have moved through all of the specified elements will be considered when updating the stream. Most of the described features available in *ProC/B* are derived from the measures that *HIT* offers, thus allowing an easy transformation from *ProC/B* to *HIT*.

Currently when analyzing the model with *HIT*, during simulation streams of data are generated, which are basically lists of pairs consisting of a time stamp and some associated value. This data is used to calculate the usual characteristics like mean, standard deviation and confidence intervals for the different measures. The *ProC/B* toolset contains a tool, that generates plots and visualizes the simulation results.

When using *OMNeT++* for simulation the key features like itemizing streams as well as the output data of the simulation should be preserved, so that this new simulation environment fits into our existing toolset. While *OMNeT++* offers basic facilities for measurement in e.g. communication protocols, they need to be extended to meet the demands for the simulation of logistics networks.

In the remainder of this section it is shown how the measurement is implemented for Standard-FUs like servers or storages. Measuring properties at composed FUs requires some additional effort and is presented afterwards. Finally the itemizing of streams is explained.

For Standard-FUs the measurement streams have to be updated when a process enters (which means a service of the FU has been requested by a process) or leaves the FU.

In the *OMNeT++* representation of the model a service request is indicated by a message sent to the FU. The population is updated whenever a process enters or leaves the FU, throughput and response time are updated when a process leaves the FU. For Standard-FUs (like server or storage) the data collection and evaluation is implemented as C++-Code within the corresponding simple modules. This brings up problems for composed FUs: When the *ProC/B* model is translated to an *OMNeT++* representation, composed FUs are represented as compound modules, thus only a NED description exists that lacks the ability to implement code for measurements. Therefore the module of every composed FU contains a specific simple module called *FUMeasures* (see Fig. 8) to realize measurements in composed FUs.

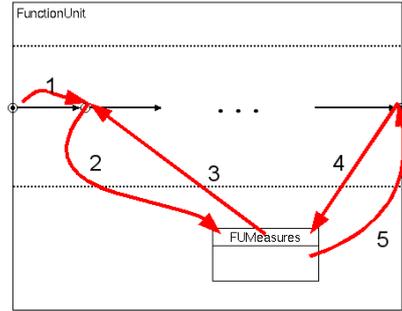


Figure 8: Message flow for the measurement in composed FUs

Fig. 8 shows the message flow that is necessary for measurements in composed FUs. Starting at the source a message is sent to the Process-ID (Process-IDs are used for the identification of a process chain and the declaration of local variables and are always connected with the source of a process chain). From there a message is sent to the special module *FUMeasures*. This module has no counterpart in *ProC/B* and its sole purpose is to enable the measurement in composed FUs. After the message is returned to the Process-ID, further elements of the process chain are processed (denoted by three dots in Fig. 8). When the process has reached the sink, messages to the module *FUMeasures* and back to the sink are sent again. All update operations of streams for a FU are performed within *FUMeasures*. The first message (sent by the Process-ID) means that a process has entered the FU (and thus the population is updated for example), the second message (sent by the sink) means that a process is leaving the FU again.

As already mentioned before, *ProC/B* allows one to specify a path consisting of elements, so that only processes, that moved along that path through the model will be considered when updating measures. Those paths are part of the *ProC/B* experiment description and need to be translated to the *OMNeT++* model and taken into account when updating the measurement streams. When mapping to *OMNeT++* the *ProC/B* experiment description is saved in an **.ini*-file, that is loaded when the simulation starts. Additionally, a parameter is set for each element appearing in one of the paths in the corresponding **.ned*-file when the model structure is mapped. During the simulation run the path a message took through the model is saved and compared with the paths, that have been specified in the *ProC/B* ex-

periment. To store this information a new message class is used, that can carry the path information. Updates of the measurement streams are only performed when the path of the message matches one of those paths from the *ProC/B* experiment.

OMNeT++ provides several classes for the collection of data and the generation of statistical measures like mean or standard deviation which are derived from the abstract class *CStatistic*. Because the calculation of means does not match the specification of the streams in *ProC/B*, our *ProC/B* to *OMNeT++* implementation provides its own classes (derived from *CStatistic*) to generate statistics (one for each of the stream types *event*, *state* and *count* mentioned before). For the estimation of confidence intervals the batch means-method [8] is used. The generated output is saved in the same format as the output of *HIT*, so that our existing tools can parse the data for result presentation.

5. VALIDATION OF THE TRANSFORMATION

If different tools are used to simulate a model, then it is necessary that the semantics of the model is the same in every simulation tool. Ideally, identity of semantics should be proved formally. Such a formal proof would require a formal semantics such that formal analysis techniques are applicable to check equivalence. Unfortunately, simulation models are much too complex to be described in simple languages that allow a formal analysis. This implies that identical behavior of models cannot be strictly verified, it can only be validated.

ProC/B was designed to introduce a well defined semantics and an automated analysis to hierarchical process chain models. A specification was laid down in [2], describing the semantics of PCEs and FUs in an operational form. Many aspects of the operational semantics are implicitly defined by the *HIT* runtime environment. Thus, the behavior of the *HIT* simulation model is the behavior which should be observed when *OMNeT++* executes the model. However, the operational semantics depends on several aspects like execution order of simultaneous events, the order of initialization which are generally not well defined in discrete event simulation and, additionally, the realization of random processes that depend on the random number generator.

We distinguish between the validation using models with and without random numbers. The former will be named deterministic models, although this is not strictly correct since simultaneous events may yield a non deterministic behavior. For the deterministic models behavior can be compared using traces. Although, *HIT* and *OMNeT++* both have a trace function it is not recommended to use these functions for comparisons since the format differs and cannot be easily transformed from one to another. Instead models are augmented with code PCEs including output statements. Such PCEs can be added to every PC. Thus, *HIT* and *OMNeT++* generate the same trace output which can be easily compared.

To prove equality of traces we developed an automated testing environment to compare output of *ProC/B* models analyzed with *HIT* and *OMNeT++*. It is based on a set of simple and deterministic *ProC/B* models, designed to test the behavior of exactly one element of *ProC/B*'s language. Driven by our batch testing environment, identically for-

mated output of *HIT* and *OMNeT++* is compared by an awk script, highlighting differences in measurement results and event traces. Additionally, a selection of deterministic models taken from former projects is also subject to comparison, making sure that our *ProC/B* language elements implemented in modules interact correctly.

Testing nondeterministic models is limited since different random number generators are used in *HIT* (actually implementations of SIMULA) and *OMNeT++*. So, even starting with same seeds, results and event orders will differ. Consequently, we can only check in a statistical sense whether the implementation is correct, i.e., the different language elements behave identically. For this purpose, animations can be compared, traces can be visualized and results can be compared using statistical test. A typical approach is to estimate the same measure with both simulators, *HIT* and *OMNeT++*, and then statistically evaluate a random variable describing the difference between both measures. This can be done by comparing confidence intervals or using statistical tests (for details see e.g. [12][chap. 10]).

Table 1 shows simulation results of an M/M/1 system ($\rho = 0.5$) as an example of simple nondeterministic system included in our testing environment. Results are sufficiently close to assume an equivalent behavior for this model with a high significance probability. We additionally compared several simulation results of larger models and obtained similar minor differences (cf. Sect. 6)

Table 1: M/M/1 system simulation results (90% confidence interval)

	Population	Throughput	Response time
<i>HIT</i>	1.00202 ±0.18%	1.00023 ±0.05%	1.001789 ±0.16%
<i>OMNeT++</i>	1.00126 ±0.33%	0.99964 ±0.14%	1.000905 ±0.22%

6. COMPARISON OF PERFORMANCE AND SIMULATION RESULTS

Though our implementation of *ProC/B* on *OMNeT++* is not as mature as the one on *HIT*, we achieved promising runtime results. Times in table 2 were taken for analyzing a model over 10.000.000 time units omitting model initialization and output.

Table 2: Runtime comparison

Model	<i>HIT</i>	<i>OMNeT++</i>
M/M/1	9min. 22sec.	3min. 30 sec.
Freight Village	7min. 01sec.	2min. 16sec.

OMNeT++ beats *HIT* by saving approximately 50 to 60% execution time.

The values given for *OMNeT++* are preliminary as we focussed on correct mapping of behavior and ignored performance issues for the time being. Performance bottlenecks still exist in statistical methods and dynamic search of matching function units to PCEs.

Table 3 shows some simulation results for the model of the freight village from section 2.1: The table contains population, throughput and response time for the server `forklifts` (see Fig. 2) estimated with `HIT` and `OMNeT++`. As one can see the results are quite similar aside from minor differences caused by different random number generators used in `HIT` and `OMNeT++`.

Table 3: Comparison of simulation results for the server forklifts (90% confidence interval)

	Population	Throughput	Response time
<code>HIT</code>	3.33275 ±0.2846%	0.36656 ±0.1257%	9.09204 ±0.2266%
<code>OMNeT++</code>	3.32575 ±0.3501%	0.36667 ±0.1251%	9.07027 ±0.3179%

7. CONCLUSIONS

Although `OMNeT++` has mainly been designed for the simulation of communication protocols, it can be used also in other areas. In this paper we demonstrated how process chain descriptions specified by `ProC/B` models can be simulated with `OMNeT++`. Since the world views of `ProC/B` and `OMNeT++` differ, the transformation is not straightforward and has to respect several special features of `ProC/B`. For example: Elements of the behavior description, like PCEs, are mapped to nodes, i.e. structural components, in the `OMNeT++` description in order to exploit `OMNeT++`'s animation capabilities. Furthermore additional elements for measurements are created, like node `FUMeasures`.

The “correctness” of the transformation has been validated by several test models where we inserted special output commands, so that discrepancies from the execution via `OMNeT++` and the reference simulator `HIT` can be detected automatically.

The current implementation is a first prototype and future work will concentrate on further improvements of the simulation efficiency and the connection to existing `ProC/B` tools. As visual aspects, positioning of `ProC/B` language elements in `OMNeT++/TKenv` will be enhanced to match `ProC/B`-editor layout and animation could include the number of requests to function units.

One of the next steps will be the utilization of Akaroa parallel simulation libraries to reduce runtimes by using multiple computers in parallel.

8. REFERENCES

- [1] F. Bause, H. Beilner, M. Fischer, P. Kemper, and M. Völker. The ProC/B Toolset for the Modelling and Analysis of Process Chains. In T. Field, P. G. Harrison, J. T. Bradley, and U. Harder, editors, *Computer Performance Evaluation / TOOLS*, volume 2324 of *Lecture Notes in Computer Science*, pages 51–70. Springer, 2002.
- [2] F. Bause, H. Beilner, and M. Schwenke. Semantik des ProC/B-Paradigmas. Technical Report 03001, ISSN 1612-1376, Sonderforschungsbereich 559 “Modellierung großer Netze in der Logistik”, 2003.
- [3] H. Beilner, J. Mäter, and N. Weißenberg. Towards a performance modelling environment: News on HIT. In R. Puigjaner and D. Potier, editors, *Modeling techniques and tools for computer performance evaluation*, pages 57–75, 1989.
- [4] H. Beilner, J. Mäter, and C. Wysocki. The Hierarchical Evaluation Tool HIT. In *Short Papers and Tool Descriptions of the 7th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, 1994.
- [5] P. Buchholz, D. Müller, P. Kemper, and A. Thümmler. OPEDo: a tool framework for modeling and optimization of stochastic models. In L. Lenzi and R. L. Cruz, editors, *VALUETOOLS*, page 61. ACM, 2006.
- [6] P. Buchholz and C. Tepper. Functional Analysis of Process Oriented Systems. In H. Fleuren, D. den Hertog, and P. Kort, editors, *Operations Research Proceedings*, pages 127–135. Springer, 2005.
- [7] Collaborative Research Center 559 “Modelling of Large Logistics Networks”. <http://www.sfb559.uni-dortmund.de>.
- [8] G.S. Fishman. *Discrete-Event Simulation Modeling, Programming and Analysis*. Springer, 2001.
- [9] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design patterns: elements of reusable object-oriented software*. Addison-Wesley Professional, 1995.
- [10] J. Huang. Simulative Bewertung von ProC/B-Modellen. Master’s thesis, Universität Dortmund, Fachbereich Informatik, Lehrstuhl 4, Dortmund, 2006.
- [11] ARIS business simulator, 2007. URL:<http://www.ids-scheer.de/>.
- [12] W. D. Kelton and A. Law. *Simulation Modeling and Analysis*. McGraw Hill, 2000.
- [13] P. Kemper and C. Tepper. Traviando - Debugging Simulation Traces with Message Sequence Charts. In *QEST*, pages 135–136. IEEE Computer Society, 2006.
- [14] Omnet++ community side. URL:<http://www.omnetpp.org/>.
- [15] J. Rathmell and D. T. Sturrock. Arena: the arena product family: enterprise modeling solutions. In Snowdon and Charnes [19], pages 165–172.
- [16] M. W. Rohrer and I. McGregor. AutoMod: simulating reality using AutoMod. In Snowdon and Charnes [19], pages 173–181.
- [17] M. D. Rossetti and H.-T. Chan. Supply chain management simulation: a prototype object-oriented supply chain simulation framework. In S. E. Chick, P. J. Sanchez, D. M. Ferrin, and D. J. Morrice, editors, *Winter Simulation Conference*, pages 1612–1620. ACM, 2003.
- [18] D. W. Schunk and B. M. Plott. Using simulation to analyze supply chains. In *Winter Simulation Conference*, pages 1095–1100, 2000.
- [19] J. L. Snowdon and J. M. Charnes, editors. *Proceedings of the 34th Winter Simulation Conference: Exploring New Frontiers, San Diego, California, USA, December 8-11, 2002*. ACM, 2002.
- [20] Q. Zhu. Beschreibung von ProC/B-Modellen zur simulativen Bewertung. Master’s thesis, Universität Dortmund, Fachbereich Informatik, Lehrstuhl 4, Dortmund, 2006.