

Applying the UML Performance Profile: Graph Grammar-based derivation of LQN models from UML specifications

Dorina C. Petriu and Hui Shen
Carleton University, Ottawa Canada, K1S 5B6

Abstract – The Object Management Group (OMG) is in the process of defining a UML Profile for Schedulability, Performance and Time that will enable the construction of models for making quantitative predictions regarding these characteristics. The paper proposes a graph-grammar based method for transforming automatically a UML model annotated with performance information into a Layered Queueing Network (LQN) performance model. The input to our transformation algorithm is an XML file that contains the UML model in XML format according to the standard XMI interface. The output is the corresponding LQN model description file, which can be read directly by existing LQN solvers. The LQN model structure is generated from the high-level software architecture and from deployment diagrams indicating the allocation of software components to hardware devices. The LQN model parameters are obtained from detailed models of key performance scenarios, represented as UML interaction or activity diagrams.

Index Terms – software performance analysis, Layered Queueing Networks (LQN), architectural patterns, Unified Modeling Language (UML), XMI, UML Performance Profile.

1. Introduction

The Object Management Group (OMG) is in the process of defining a UML Profile for Schedulability, Performance and Time [9] that would enable the construction of models that can be used for making quantitative predictions regarding these characteristics. The original RFP issued in March 1999 was followed by the first “Response to RFP” submission in August 2000, and by a revised submission in June 2001 [9]. The later includes some additional aspects not covered in the former, among which is a

section dealing with performance analysis. The proposed performance profile extends the UML metamodel with stereotypes, tagged values and constraints, which make possible to attach performance annotations (such as resource demands and visit ratios) to a UML model. In order to conduct quantitative performance analysis of an annotated UML model, one must first translate it into a performance model, use an existing performance analysis tool for solving the performance model and then import the performance analysis results back in the UML model. The focus of this paper is the first step of the process. The paper proposes a graph-grammar based method for transforming a UML model annotated with performance information into a Layered Queueing Network (LQN) performance model [18, 19]. The input to our transformation algorithm is an XML file that contains an annotated UML model in XML format according to the standard XMI interface [8], and the output is the corresponding LQN model description file, which can be read directly by existing LQN solvers [4]. The proposed transformation algorithm was implemented in Java. More work is necessary for the present implementation to become a tool prototype, especially the addition of a suitable GUI.

The transformation approach is as follows: the LQN model structure is generated from the high-level software architecture, more exactly from the architectural patterns showing high-level software components and their relationships, and from deployment diagrams indicating the allocation of software components to hardware devices. The LQN model parameters are obtained from detailed models of key performance scenarios, represented as UML interaction or activity diagrams, annotated with performance information. The paper continues the previous work of the authors from [10, 11, 12].

Since the introduction of the Software Performance Engineering technique in [17], there has been a significant effort to integrate performance analysis into the software development process throughout all lifecycle phases. This requires the ability to derive performance models from software design specifications. A survey of the techniques developed in the recent years for deriving performance models from UML models is given in [1]. Among the technique surveyed, those from [2] and [7] are

the closest to our approach. However, our paper is a unique combination of the following characteristics: a) it accepts as input XML files produced by UML tools, b) it generates LQN performance models by applying graph transformation techniques to graphs of metaobjects that represent different UML diagrams of the input model, and c) it uses the UML performance profile for adding performance annotations to the input model.

2. Background: UML Performance Profile, LQN model and Architectural Patterns

2.1 UML Performance Profile

According to [9], the UML Performance Profile provides facilities for:

- capturing performance requirements within the design context
- associating performance-related QoS characteristics with selected elements of the UML model
- specifying execution parameters which can be used by modelling tools to compute predicted performance characteristics
- presenting performance results computed by modelling tools or found by measurement.

The Profile describes a domain model, shown in Fig. 1, which identifies basic abstractions used in performance analysis. *Scenarios* define response paths through the system, and can have QoS requirements such as response times or throughputs. Each scenario is executed by a job class, called here a *workload*, which can be closed or open and has the usual characteristics (number of clients or arrival rate, etc.) Scenarios are composed of *scenario steps* that can be joined in sequence, loops, branches, fork/joins, etc. A scenario step may be an elementary operation at the lowest level of granularity, or may be a complex sub-scenario composed of many basic steps. Each step has a mean number of repetitions, a host execution demand, other demand to resources and its own QoS characteristics. Resources are another basic abstraction, and can be active or passive, each with their own attributes. The Performance profiles maps the classes from Fig. 1 to a stereotype that can be applied to a number of UML model elements, and each class attribute to a tagged value. For example,

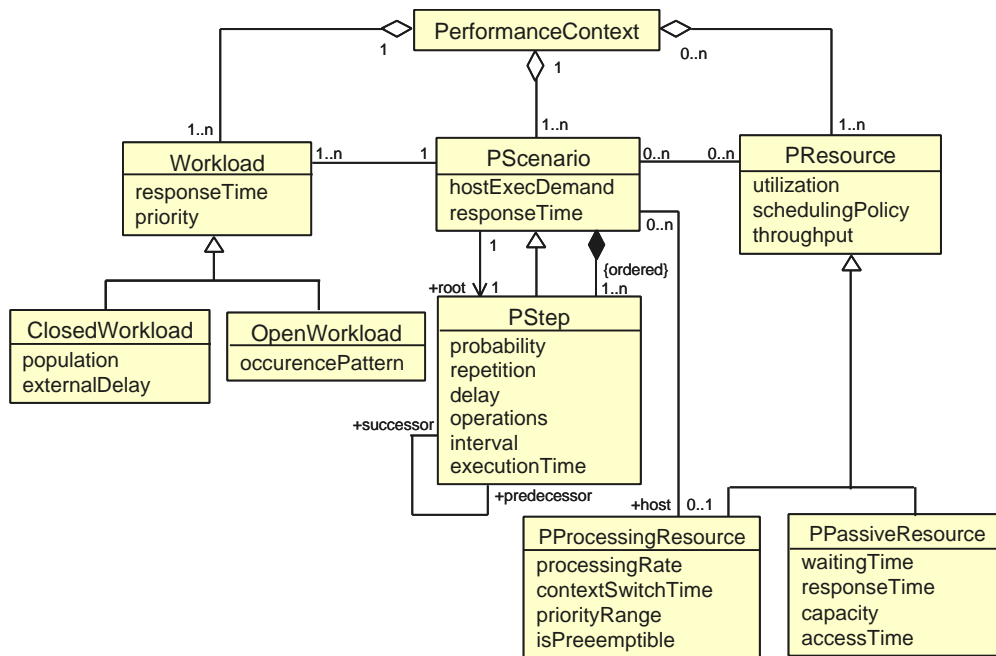


Figure 1 Domain model in the UML Performance Profile

the basic abstraction PStep is mapped to the stereotype <<Pastep>> that can be applied to the following UML model elements: Message and Stimulus (when the scenario is represented by an interaction diagram) or ActionState and SubactivityState (when the scenario is represented by an activity diagram).

In our implementation, we process XML files produced by current UML tools, which obviously do not support the Performance Profile yet. Therefore, we have attached the tagged values associated with the stereotypes “by hand” to different model elements.

Fig.2 illustrates the inter-operability of the different tools involved: a UML tool (such as rational Rose or ArgoUML, a free source tool), a performance model solver (LQN analytical solver or simulator) and our UML to LQN transformation implementation. So far we have made progress on the forward path (represented with black arrows) but have not attempted yet the backward path (in gray arrows).

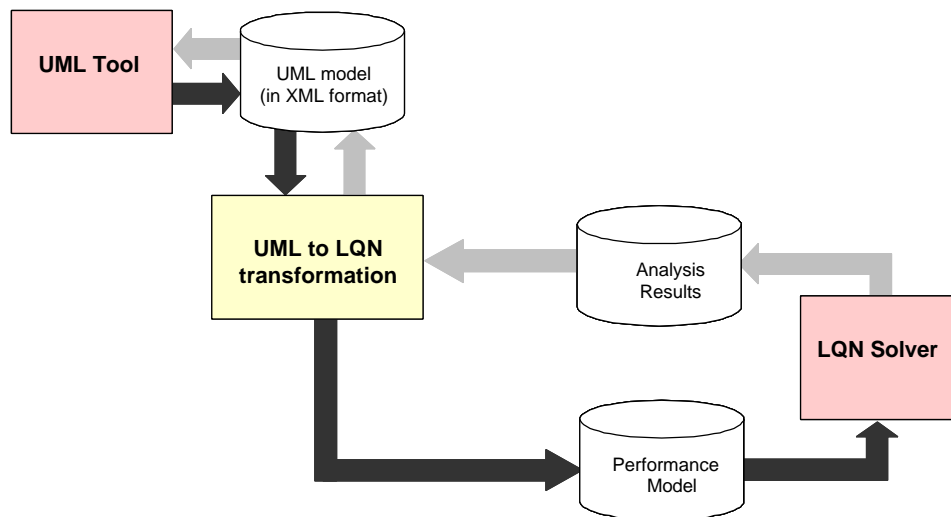


Figure 2 Tool inter-operability

2.2. The LQN model

LQN was developed as an extension of the well-known QN model, at first independently in [18, 19] and [14], then as a joint effort [4]. The LQN toolset presented in [4] includes both simulation and analytical solvers. The main difference with respect to QN is that LQN can easily represent nested services: a server which receives and serves client requests, may become in turn a client to other servers from which it requires nested services while serving its own clients. A LQN model is represented as an acyclic graph, whose nodes represent software entities and hardware devices, and arcs denote service requests. The software entities (also known as *tasks*) are drawn as parallelograms, and the hardware devices as circles. The nodes with outgoing but no incoming arcs play the role of *clients*, the intermediate nodes with both incoming and outgoing arcs are usually software servers and the leaf nodes are hardware servers (such as processors, I/O devices, communication network, etc.) A software or hardware server node can be either a single-server or a multi-server. Fig. 1 shows a simple example of an LQN model of a web server: at the top there is a customer class with a given number of stochastic identical clients. Each client sends demands for different services of the WebServer. Each kind of service offered by an LQN task is modelled as a so-called *entry*, drawn as a parallelogram “slice” in the figure. Every entry has its own execution times and demands for other services (given as

model parameters). In this case, the WebServer entries require services from different entries of the Database task. Each software task is running on a processor shown as a circle. Also as circles are shown the communication network delays and the disk device used by the Database. It is worth mentioning that the word *layered* in the LQN name does not imply a strict layering of tasks (for example, tasks in a layer may call each other or skip over layers). All the arcs used in this example represent *synchronous* requests, where the sender of a request message is blocked

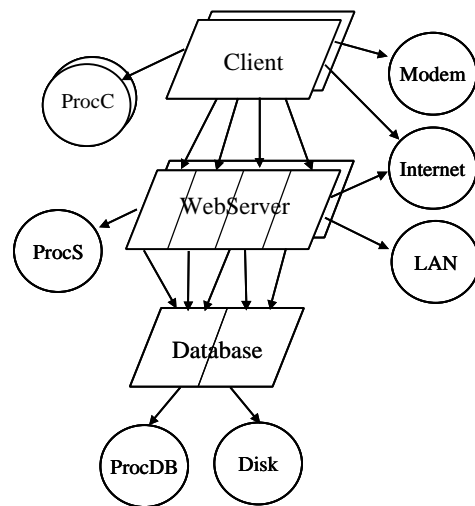


Figure 3 Simple LQN model

until it receives a reply from the provider of service. It is possible to have also *asynchronous* request messages, where the sender does not block after sending a request and the server does not send any reply back. Another communication style in LQN allows for a client request to be processed by a chain of servers instead of a single server, as shown in section 4. The first server in the chain will *forward* the request to the second, etc., and the last server in the chain will *reply* to the client. Although not explicitly illustrated in the LQN notation, every server, be it software or hardware, has an implicit message queue where incoming requests are waiting their turn to be served. Servers with more than one entry have a single input queue, where requests for different entries wait together.

A server entry may be decomposed in two or more sequential *phases* of service. *Phase 1* is the portion of service during which the client is blocked, waiting for a reply from the server (it is assumed that the client has made a synchronous request). At the end of phase 1, the server will reply to the client, which will unblock and continue its execution. The remaining phases, if any, will be executed in parallel with the client. A recent extension to LQN [5] allows for an entry to be further decomposed into *activities* if more details are required to describe its execution. The activities are connected together to form a directed graph, which may branch into parallel threads of control, or may choose randomly between

different branches. Just like phases, activities have execution time demands, and can make service requests to other tasks. An example of LQN with activities is given in section 4, Fig. 8.

The parameters of a LQN model are as follows:

- customer (client) classes and their associated populations or arrival rates;
- for each phase (activity) of a software task entry: average execution time;
- for each phase (activity) making a request to a device: average service time at the device, and average number of visits;
- for each phase (activity) making a request to another task entry: average number of visits
- for each request arc: average communication delay;
- for each software and hardware server: scheduling discipline.

2.3. Architectural Patterns

In our approach, the structure of the LQN model is generated from the high-level architecture, and more exactly from the architectural patterns used in the system. Frequently used architectural solutions are identified in literature as *architectural patterns* (such as pipeline and filters, client/server, client/broker/server, layers, master-slave, blackboard, etc.) A pattern introduces a higher-level of abstraction design artifact by describing a specific type of collaboration between a set of prototypical components playing well-defined roles, and helps our understanding of complex systems. Each architectural pattern describes two inter-related aspects: its *structure* (what are the components) and *behaviour* (how they interact). In the case of high-level architectural patterns, the components are usually concurrent entities that execute in different threads of control, compete for resources, and their interaction may require some synchronization. The patterns are represented as UML *collaborations* (not to be confused with UML *collaboration diagrams*, a type of interaction diagrams) [9]. The symbol for a collaboration is an ellipse with dashed lines that may have an “embedded” square showing the roles played by different pattern participants.

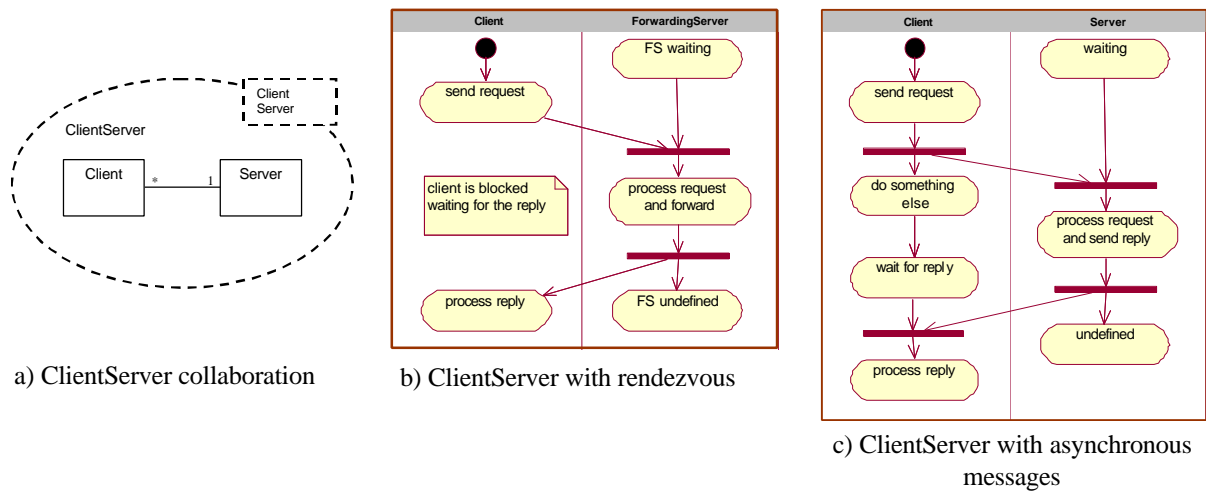


Figure 4. ClientServer architectural pattern

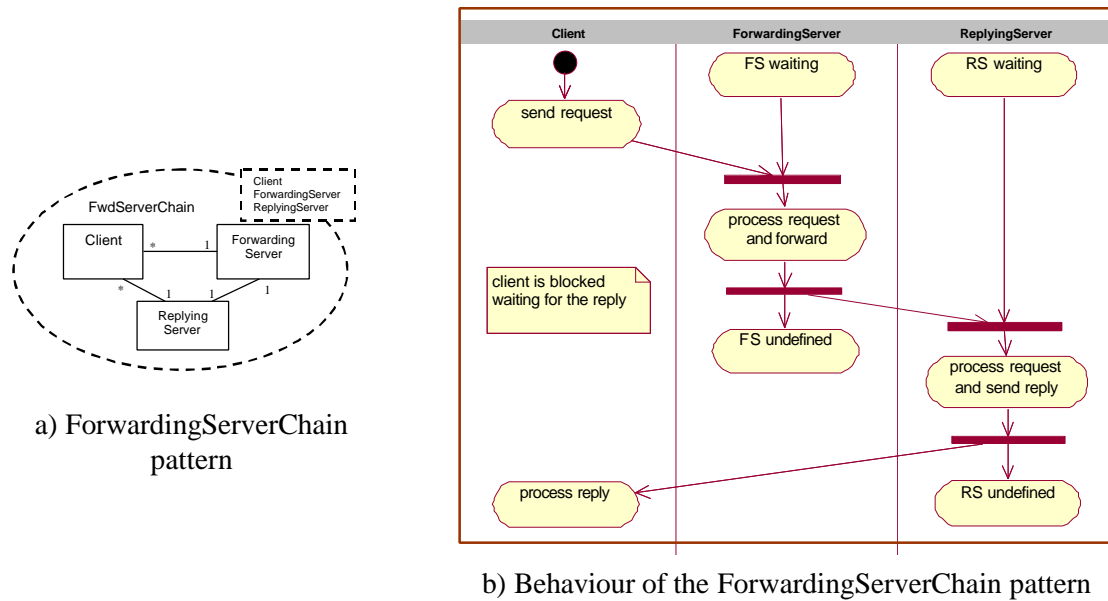


Figure 5. ForwardingServerChain architectural pattern

In Fig. 4 and 5 are shown the structure and behaviour of two patterns used in our case study: ClientServer and ForwardingServerChain. The ClientServer pattern has two alternatives: the one shown in Fig.4.b is using a rendezvous communication style (where the client sends the requests then remains blocked until the sender replies), whereas the one from Fig. 4.c is using an asynchronous communication style (where the client continues its work after sending the request, and later on will

accept the server's replay). The ForwardingServerChain, shown in Fig.5, is an extension of the ClientServer pattern, where the client's request is served by a series of servers instead of a single one.

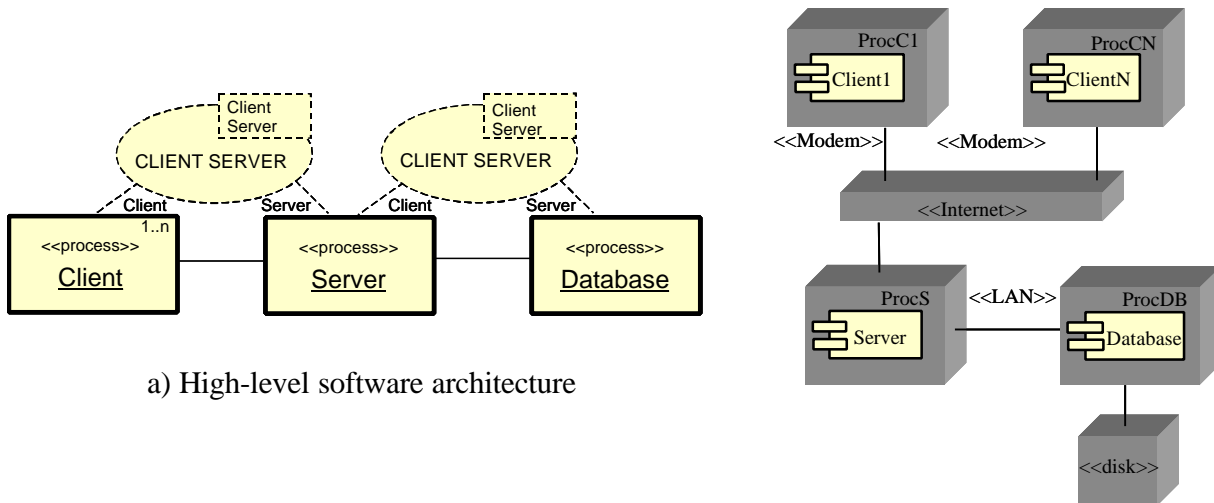
There may be more than two servers in the chain (only two are shown in Fig.5). The servers in the middle play the role of ForwardingServer, as each one forwards the request to the next server in the chain after doing their part of service. The last server in the chain plays the role of ReplyingServer as it sends the reply back to the client. In this paper we show how these two patterns are converted into LQN models. More architectural patterns and the corresponding rules for translating them into LQN are described by the authors of the present paper in [10, 11].

3. Transformation from UML to LQN

Similar to the SPE methodology from [16, 17], the starting point for our algorithm is a set of key performance scenarios annotated with performance information. For each scenario we derive an LQN submodel, then merge the submodels together. The approach for merging is similar with the one used in [6], where LQN submodels were derived from execution traces.

The derivation of each LQN submodel is done in two big steps:

- a) The submodel structure (i.e., the software and hardware tasks and their connecting arcs) is obtained from the high-level architecture of the UML model and from the deployment of software components to hardware devices. Two kinds of UML diagrams are taken into account in this step: a high-level collaboration diagram that shows the concurrent/distributed high-level component instances and the patterns in which they participate, and the deployment diagram. Fig.6 shows these two diagrams for the web server model that was used to derive the LQN model from Fig.3 without task entries. Each high-level software component is mapped to a LQN task, and each hardware device (processor, disk, communication network, etc.) is mapped to a LQN hardware task. The arcs between LQN nodes correspond to the links from the UML diagrams. It is important to mention that in the first transformation step from UML to LQN we take into account only the



a) High-level software architecture

Figure 6. UML diagrams used to generate the structure of the LQN model from Figure 3

structural aspect of the architectural patterns; their behavioural aspect will be considered in the next step.

- b) LQN task details are obtained from UML scenario models represented as activity diagrams, over which we overlay the behavioural aspect of the architectural pattern, making sure that the scenario is consistent with the patterns. By “LQN details” we understand the following elements of each task: entries, phases, activities (if any) and their execution time demands and visit ratio parameters, as described in section 2.2. A task entry is generated for each kind of service offered by the corresponding software component instance. The services offered by each instance are identified by looking at the messages received by it in every scenario taken into account for performance analysis.

Scenarios can be represented in UML by sequence, collaboration or activity diagrams. (The first two are very close as descriptive power and have similar metamodel representation). UML statecharts are another kind of diagrams for behaviour description, but are not appropriate for describing scenarios. A statechart describes the behaviour of an object, not the cooperation between several objects, as needed in a scenario.

In the proposed approach, we decided to use activity diagrams for translation to LQN. The main reason is that sequence (collaboration) diagrams are not well defined in UML yet, as they are lacking convenient features for representing loops, branches and fork/join structures. Other authors who are building performance models from UML designs have pointed out this deficiency of the present UML standard, and are using instead extended sequence diagrams that look like the Message Sequence Chart standard (see [16], [17] for well known examples). We did not take the approach of extending the sequence diagrams with the missing features because our algorithm takes in XML files generated by UML tools, and parses graphs of UML metaobjects. Our implementation is consistent with the present UML metamodel and XMI interface; moreover, it uses a free-source library that implements the UML metamodel as defined in [8]. Therefore, our choice was to use activity diagrams, which are able to represent branch/merge, fork/join and activity composition.

However, the activity diagrams have a disadvantage with respect to sequence (collaboration) diagrams: they do not show what objects are responsible for different actions. An attempt to counterbalance this weakness was the introduction of “swimlanes” (or partitions) in the UML standard. A swimlane contains actions that are performed by a certain instance or set of instances (for example, a swimlane can be associated to a whole department when modeling a workflow problem). In our approach, we associate a swimlane with each concurrent (distributed) component, which will be translated into a LQN task (see Fig. 5, 6 and X). Since many UML modellers prefer sequence (collaboration) diagrams for expressing the cooperation between objects, we proposed in [12] an algorithm based on graph transformations for converting automatically sequence diagrams into activity diagrams. The transformation associates a swimlane to all the objects that belong to a concurrent (distributed) component, and therefore adjusts the level of abstraction of the model to our needs.

4. From Activity Diagrams to LQN entries, phases and activities

This section presents the graph grammar-based transformation of activity diagrams into LQN detailed features (i.e., the realization of step (b) from the previous section). The graph-grammar formalism is appropriate in this case because both UML and LQN models are described by graphs. The essential idea of all *graph grammars* or *graph rewriting systems* is that they are generalization of the *string grammars* that are used in compilers. The terms “graph grammars” and “graph rewriting systems” are often considered synonymous. However, a graph grammar is a set of production rules that generates a language of terminal graphs and produces nonterminal graphs as intermediate results, whereas a graph rewriting system is a set of rules that transforms one instance of a given class of graphs into another instance of the same class of graphs, without distinguishing between terminals and nonterminals graphs. The main component of a graph grammar is a finite set of production rules. A production is a triple (L,R,E) , where L and R are graphs (the left-hand side and right-hand side, respectively) and E is some embedding mechanism. Such a production rule can be applied to a host graph H as follows: when an occurrence of L is found in H , it is removed and replaced with a copy of R ; finally, the embedding mechanism E is applied to attach R to the remainder of H . [13]

In our case, the initial host graph is the set of metaobjects that represents a given activity diagram (the metaobjects are the nodes and the links between them are the edges of the graph). According to the UML metamodel, the nodes (i.e., metaobjects) are of type *StateVertex* and *Transition*. A *StateVertex* type has a number of subtypes, among which *State* and *Pseudostate* are the most important. *State* is associated eventually with the actions represented in the diagram, whereas the special diagram blocks such as “fork”, “join”, “choice”, etc., are *Pseudostates* (see [8] for more details).

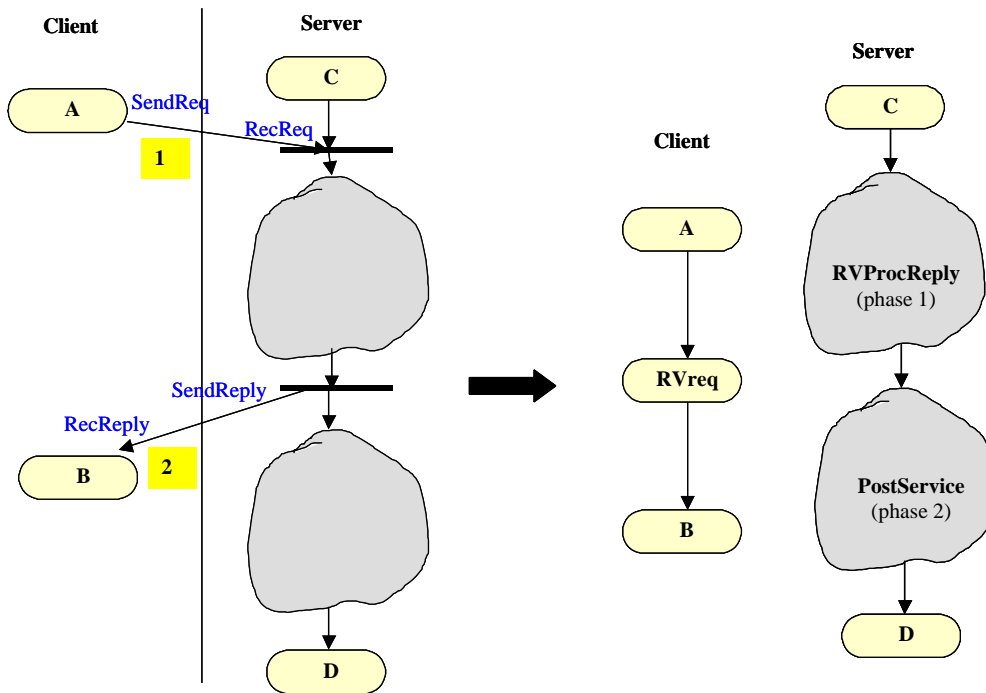
Our purpose is to parse the activity diagram graph to check first whether it is correct, then to divide it into subgraphs that correspond to various LQN elements (entries, phases, etc.). In general, parsing graph grammars is quite difficult, and in some cases even impossible [13]. In this case, we have found

a shortcut, as explained below, that decomposes the original host graph into subgraphs described by simpler graph-grammars, very similar to string grammars. Each subgraph corresponds to a partition and describes the behaviour of a single component, dealing with sequences of scenario steps, loops, alternative branches and fork/join structures. We were able to define and implement a top-down parser with recursive methods to parse these subgraphs. Our algorithm for step (b) contains two substeps:

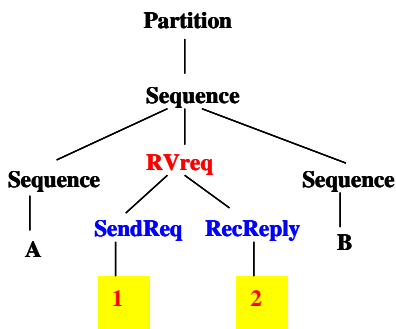
b.1) Overlay the behaviour of the architectural patterns extracted in step (a) from the high-level collaboration diagram over the activity diagram, and verify whether the communication between concurrent components is consistent with the pattern. This is done by traversing the graph, by identifying the cross-transitions between swimlanes, and by checking if they follow the protocol defined by the pattern. Then, apply a graph transformation rule corresponding to the respective pattern, and attach appropriate nonterminal symbols to the subgraphs identified (see Fig. 7 and 8). This disconnects practically the subgraphs corresponding to each swimlane from its neighbors. Repeat this step for all inter-component communication (we assume that all are covered by some architectural pattern).

b.2) Parse separately the subgraph within each swimlane. This will deal with sequences, loops, alternative branches and fork/join structures. When parsing a subgraph corresponding to a certain LQN element (phase or activity) identified in the previous step, compute its execution time S from the execution times of the contained scenario steps as follows: $S = \sum_{i=1,n} r_i S_i$, where r_i is the number of repetitions and S_i the host execution time of scenario step i .

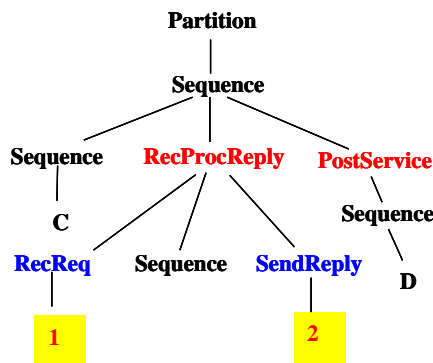
The graph transformation rule and parsing tree for a ClientServer pattern with rendezvous are illustrated in Fig. 7, whereas those for a ClientServer pattern with asynchronous communication are illustrated in Fig. 8. According to Fig.7.a, two subgraphs are identified on the server side. One corresponds to the phase 1 of service and is found between the “join” Pseudostate marking the




a) Graph transformation rule for the Client Server pattern with rendezvous communication



b) Parse tree for the Client's partition



c) Parse tree for the Server's partition




d) Generated LQN

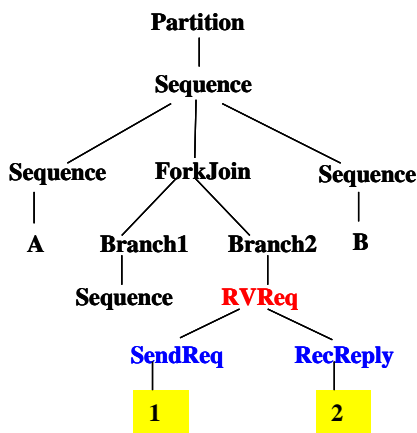
Figure 7. Transformation rule and parsing in the case of a ClientServer pattern with rendezvous communication

receiving of the client request and the “fork” Pseudostate marking the sending of the reply. The other subgraph corresponds to the second phase of service, and is found between the “fork” marking the sending of the reply and either a “waiting” state for a new request or the “undefined” state we use by default to mark the end of the respective component behaviour on behalf of the current scenario. (This subgraph may very well be empty). The subgraphs are labeled with the nonterminal symbols *RVProcReply* and *PostService*, respectively. After step (b.1), these subgraphs will be parsed

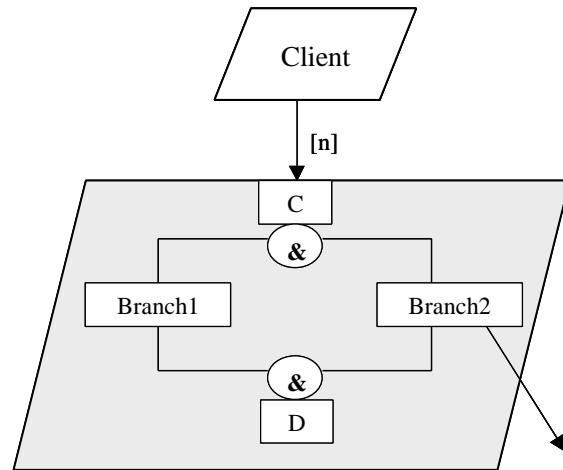
individually. Their partial parsing trees are shown in Fig. 7.b and 7.c. The LQN elements generated are shown in Fig. 7.d. A new entry of the Server task is generated for each new type of request accepted by the server. In this case, the entry has two phases, and their service time is computed as explained above. A LQN request arc is generated between the Clients phase (activity) that contains the new



a) Graph transformation rule for the ClientServer pattern with asynchronous communication



b) Parse tree for the Client's partition



c) Generated LQN elements

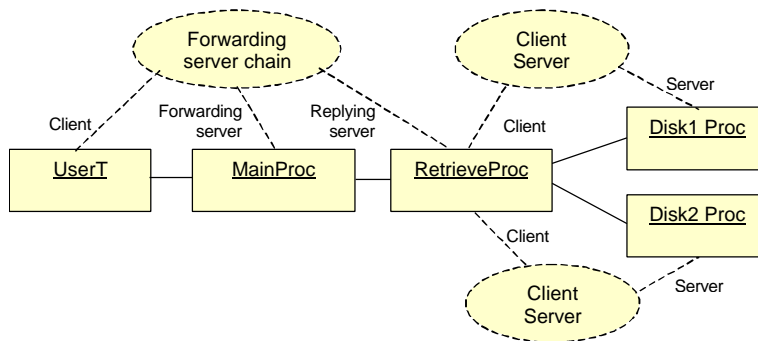
Figure 8. Transformation rule and parsing in the case of a ClientServer pattern with asynchronous communication

RVreq node inserted by the rule from Fig 7.1, whose visit ratio n is identical with the number of repetitions of the state A originating the rendezvous request.

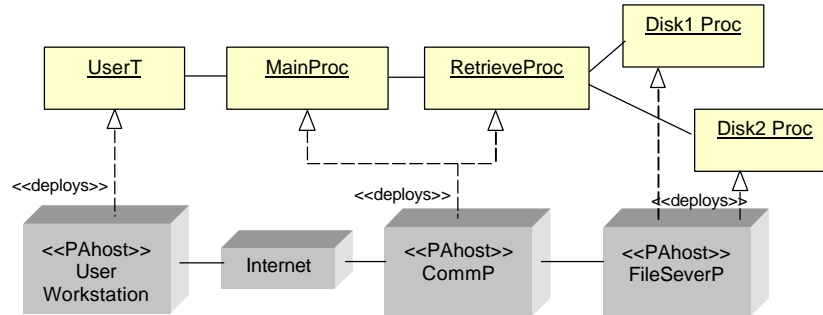
Fig.8.a shows a similar transformation rule for the ClientServer pattern with asynchronous communication. The difference here is that the client does not block immediately after sending the reply to the server. A fork/join structure is inserted in the client side, as seen in Fig. 8.a and 8.b. The LQN elements generated are shown in Fig. 8.c. The new entry contains LQN activities and a similar fork/join structure. The transformation rules for other patterns are not given here due to space limitations.

5. Case study: group communication server

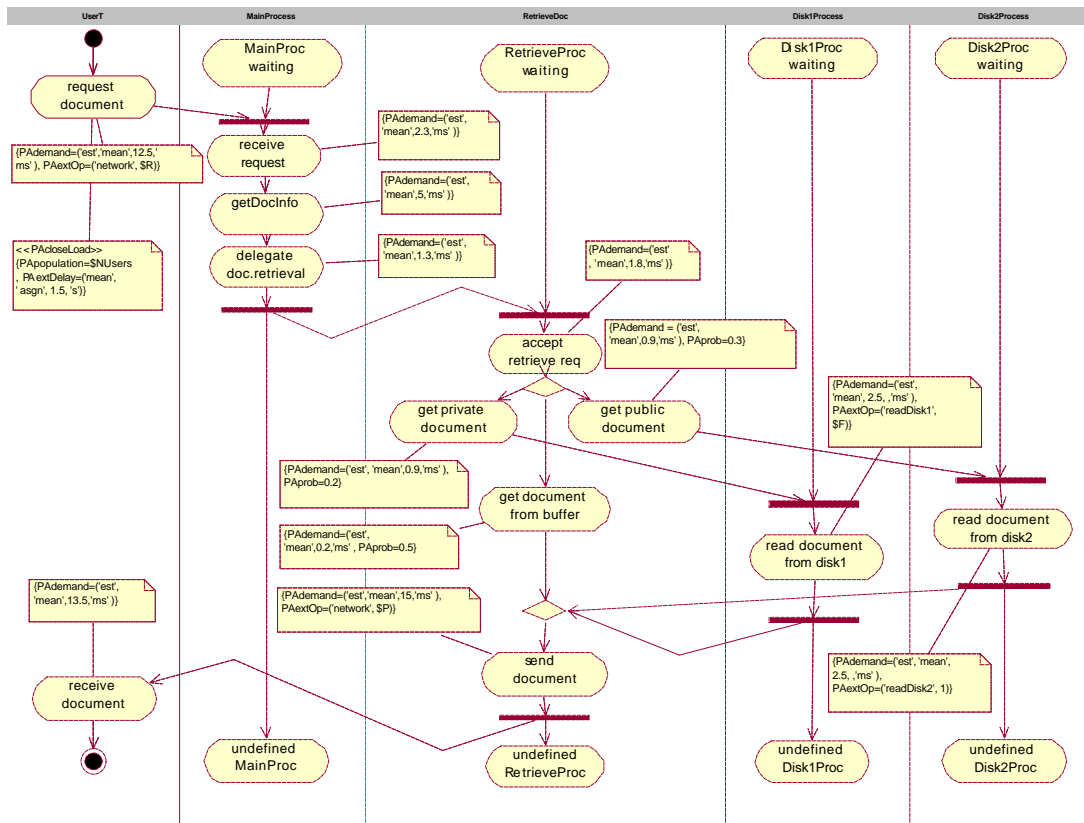
In this section is presented the application of the proposed UML to LQN transformation algorithm to a case study [15]. A group communication server accepts two classes of documents from its subscribers, private and public, each with its own access rights. The documents are kept in two different files: the private documents on disk1 and the public ones on disk2. Five use cases were analyzed for the system: subscribe, unsubscribe, submit document, retrieve document and update document. However, only one scenario, namely Retrieve, is presented here. Fig. 9.a shows the architectural patterns in which are participating the components involved in this scenario, and Fig. 9.b gives the deployment diagram. The activity diagram with performance annotations describing the scenario is given in Fig.9.c. A User process sends a request for a specific document to the Main process of the server, which determines the type of request and forwards it to another process named RetrieveProc. This process is responsible for retrieving the documents. Half of all the requests will refer to documents that will be found in the buffer, so no disk access is necessary. For the other half, 20% will be requests for private and 30% for public documents. In each case, RetrieveProc delegates the responsibility of reading the document to the corresponding disk process.



a) High-level architecture for the “Retrieve” scenario



b) Deployment diagram for the communication server



c) Activity Diagram annotated with performance info for the Retrieve scenario

Figure 9. Group communication server: the high-level architecture, deployment and activity diagram for the Retrieve scenario

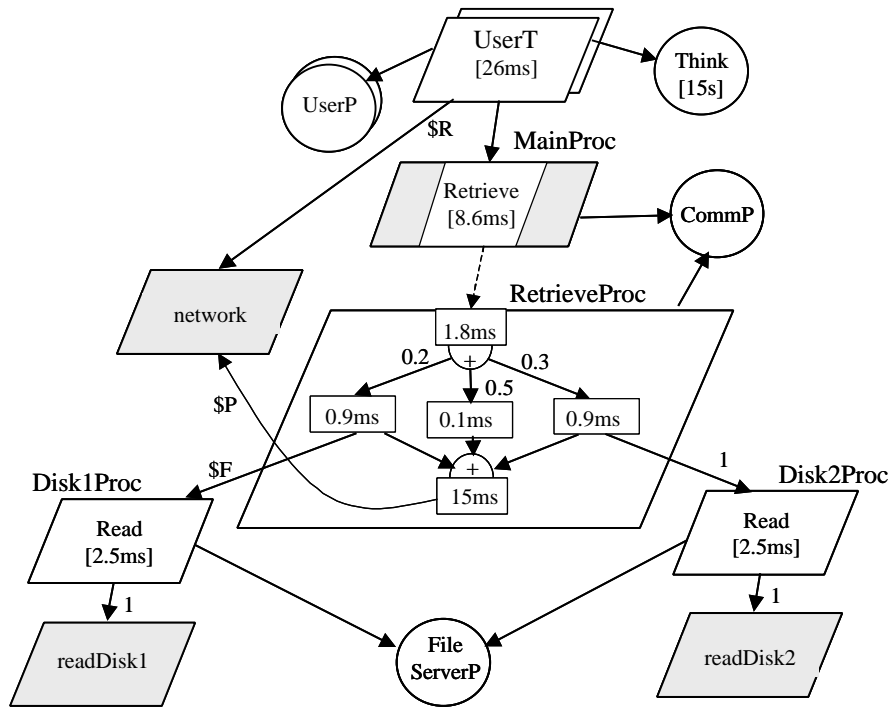


Figure 9. LQN submodel derived from the Receive scenario

Disk1Proc, will read the private documents from a sequential file, so it will make a number \$F of accesses to an external operation named “readDisk1” in order to find the desired document. On the other hand, Disk2Proc will read the public documents from an indexed file, so it will make only one access to the external operation “readDisk2”. The performance characteristics and the resources required for the external operations are described elsewhere. The external operations were specially provided in the UML performance profile to allow the UML modeler to describe the system at the right level of abstraction.

After getting the document either from file or from memory, RetrieveProc sends it back to the user. The scenario steps that send messages over the Internet invoke an external operation “network” once for every packet transmitted. This will allow to plug in directly into the performance model communication network delays that are not fully represented in the UML model. Every scenario step in Fig.9.c has a *PAdemand* tagged value indicating its estimated mean execution time on the host processor. The workload for the Retrieve scenario is closed, with a number \$Nusers of clients. A user “

“thinks” for a mean delay of 15s. (The identifiers starting with ‘\$’ indicate variables that must be assigned concrete before doing the actual performance analysis).

The LQN model obtained by applying our method is shown in Fig.10. A LQN task was generated for each of the five software components from Fig.9. Additional tasks (shown shaded in the figure) were generated for the external operations. The task MainProc has many entries, one for each type of requests received. However, only one of its entries is determined from this scenario. This entry forwards the client request to RetrieveProc (represented by a dotted request arc). The only entry of RetrieveProc has internal branching, represented by an LQN activity graph that mirrors the scenario steps from the activity diagram. The purpose of this paper is to present the proposed UML to LQN transformation, so no performance analysis results are presented here.

Our experience with the UML Performance Profile shows that it is relatively easy to understand, and that it provides enough performance annotations for generating working LQN models. However, we would like to add a few items on the “wish list” for the Performance Profile. Firstly, an additional tagged value is needed for expressing the size of messages, which is necessary for calculating the network delays. Secondly, it would be very useful to use parameterized expressions instead of concrete numbers for the tagged values. Thirdly, sometime it’s necessary to define a workload over a set of scenarios, and to give the probability of choosing each scenario in the set.

6. Conclusions

Our experience with the graph-grammar formalism shows that it is very powerful and modularized by nature. The rule-based transformation approach lends itself rather easily to extensions. We are working right now on adding new transformation rules for other architectural patterns, such as pipeline and filters, critical section, blackboard, master-slave, etc. Another kind of extension we are planning on doing is the addition of a suitable GUI. An area that is completely uncovered is the backward path to

bring back into the UML model the results from the performance analysis solver (shown in gray in Fig.2).

Regarding the inter-operability with UML tools, we have met with some problems due to the fact that the present UML tools do not support entirely the current UML standard. For example, Rational Rose does not support yet the following features: collaborations (i.e., the dashed ellipsis symbol), object flow in activity diagrams and tagged values. Another example: a free-source UML tool we have been using, named ArgoUML, does not support swimlanes and object flow in activity diagrams, etc. In order to test our algorithm, we have obtained XML files from the existing UML tools, but had to change them by hand in order to add the missing features. We hope that this problem will disappear with time, so that tool inter-operability will become a true reality.

References

- [1] S. Balsamo, M. Simeoni, "On transforming UML models into performance models", *Workshop on Transformations in the Unified Modeling Language*, Genova, Italy, April 2001.
- [2] V. Cortellessa, R. Mirandola, "Deriving a Queueing Network based Performance Model from UML Diagrams", *Proceedings of the Second International Workshop on Software and Performance*, Ottawa, Canada, pp 58-70, Sept. 2000.
- [3] H.Gomaa, D.A.Menasce, "Design and Performance Modeling of Component Interconnections Patterns for Distributed software architectures, *Proceedings of 2nd ACM Workshop on Software and Performance, WOSP'2000*, Ottawa, Canada, Sept.2000.
- [4] G. Franks, A. Hubbard, S. Majumdar, D.C. Petriu, J. Rolia, C.M. Woodside, "A toolset for Performance Engineering and Software Design of Client-Server Systems", *Performance Evaluation*, Vol. 24, Nb. 1-2, pp 117-135, November 1995.
- [5] Greg Franks, "Performance Analysis of Distributed Server Systems", Report OCIEE-00-01, Ph.D. Thesis, Carleton University, Ottawa Canada, Jan. 2000.
- [6] C.E. Hrischuk, C.M. Woodside, J.A. Rolia, "Trace-Based Load Characterization for Generating Software Performance Models, *IEEE Trans. on Software Eng.*, v 25, n 1, pp 122-135, Jan. 1999.

- [7] P.Kähkipuro, "UML-Based Performance Modeling Framework for Component-Based Distributed Systems", in R.Dumke et al.(Eds): *Performance Engineering, LNCS 2047*, Springer, pp.167-184, 2001.
- [8] Object Management Group, UML Specification Version 1.3, OMG Doc. ad/99-06-08, 1999.
- [9] Object Management Group, UML Profile for Scheduling, performance and Time, OMG Document ad/2001-06-14, <http://www.omg.org/cgi-bin/doc?ad/2001-06-14>, June 2001.
- [10] D.C.Petriu, X.Wang, "From UML description of high-level software architecture to LQN performance models", in *Applications of Graph Transformations with Industrial Relevance AGTIVE'99* (eds. M.Nagl, A. Schuerr, M. Muench), LNCS 1779, pp. 47-62, Springer, 2000.
- [11] D.C.Petriu, C.Shousha, A.Jalnapurkar, "Architecture-Based Performance Analysis Applied to a Telecommunication System", *I.E.E.E. Transactions on Software Eng.*, Vol.26, No.11, pp. 1049-1065, Nov. 2000.
- [12] D.C.Petriu, X.Wang "Consistent Behaviour Representation in Activity and Sequence Diagrams", in *UML'2000 The Unified Modeling Language - Advancing the Standard*, LNCS 1939, pp.369-382, Springer, 2000.
- [13] G.Rozenberg (ed), *Hanbook of Graph Grammars and Computing by Graph Transformation*, World Scientific, 1997.
- [14] J.A. Rolia, K.C. Sevcik, "The Method of Layers", *IEEE Trans. On Software Engineering*, Vol. 21, Nb. 8, pp 689-700, August 1995.
- [15] W.C.Scratchley, "Evaluation and diagnosis of Concurrency Architectures", Ph.D Thesis, Carleton University, Dept. of Systems and Computer Eng., 2000.
- [16] C.U. Smith, *Performance Engineering of Software Systems*, Addison Wesley, 1990.
- [17] C.U. Smith, L.G Williams, *Performance Solutions: A Practical Guide to Creating responsive, Scalable Software*, Addison Wesley, 2001.
- [18] C.M. Woodside. "Throughput Calculation for Basic Stochastic Rendezvous Networks". *Performance Evaluation*, vol.9 (2), pp. 143-160, April 1988.
- [19] C.M. Woodside, J.E. Neilson, D.C. Petriu, S. Majumdar, "The Stochastic Rendezvous Network Model for Performance of Synchronous Client-Server-like Distributed Software", *IEEE Transactions on Computers*, Vol.44, Nb.1, pp 20-34, January 1995.