



Repositorio Institucional de la Universidad Autónoma de Madrid

<https://repositorio.uam.es>

Esta es la **versión de autor** de la comunicación de congreso publicada en:
This is an **author produced version** of a paper published in:

Web Information Systems and Technologies: Third International Conference,
WEBIST 2007, Barcelona, Spain, March 3-6, 2007, Revised Selected Papers.
Lecture Notes in Business Information Processing. Springer 2008. 169-183

DOI: http://dx.doi.org/10.1007/978-3-540-68262-2_13

Copyright: © 2008 Springer-Verlag

El acceso a la versión del editor puede requerir la suscripción del recurso
Access to the published version may require subscription

Model Driven Formal Development of Digital Libraries

Esther Guerra¹, Juan de Lara², and Alessio Malizia¹

¹ Dep. Computer Science, Universidad Carlos III, Madrid (Spain),
eguerra@inf.uc3m.es, amalizia@inf.uc3m.es

² Dep. Computer Science, Universidad Autónoma, Madrid (Spain),
jdelara@uam.es

Abstract. This paper shows our model-driven approach for the formal construction and validation of Digital Libraries (DLs). We have defined a Domain Specific Visual Language (DSVL) called VisMODLE, which allows the description of a DL using five different viewpoints: services, behaviour, collections, structure and society. From a meta-model based description of the different viewpoints, we have generated a modelling environment for VisMODLE. We have provided the environment with a code generator that produces XUL code for the DL's user interface and composes the application using predefined components that implement the different services. Moreover, we have also added validation and simulation capabilities to the environment. Using the behavioural models (state-machine based), we can visually animate the system. In addition, the combined behaviour of actors and services can be transformed into a Petri net for further analysis.

Key words: Digital Libraries, Model Driven Development, Formal Methods, Meta-modelling, Graph Transformation

1 Introduction

The concept of Digital Library (DL) seems hard to be completely understood and evades definitional consensus. Licklider [10] visualized a collection of digital versions of the worldwide corpus of published literature and its availability through interconnected computers. More recently, a Delphi study [9] of DLs coalesced a broad definition: organized collection of resources, mechanisms for browsing and searching, distributed networked environments, and sets of services objectified to meet users' needs. The Presidents Information Technology Advisory Committee (PITAC) Panel on DLs treats them as the networked collections of digital texts, documents, images, sounds, scientific data and software, that make up the core of today's Internet and tomorrow's universally accessible digital repositories of human knowledge [14]. Underlying these definitions there is the consensus agreement that DLs are fundamentally complex due to its inherently interdisciplinary nature. They are usually built from scratch using specialized architectures that do not benefit from previous DLs and software design experiences. The lack of formal models leads to branching efforts and has made interoperability (in both metadata and software levels) a crucial problem in the DL field.

From the beginning of computer science, software engineers have sought methods to increase the productivity and quality of applications. A means to achieve such goal

is to increase the level of abstraction of system descriptions. In Model-Driven Software Development (MDS) [15], models are the primary asset, from which the code of the application is generated. The idea is to capitalize the knowledge in a certain application domain by providing developers with DSLs describing the domain concepts. DSLs are less error-prone than other general-purpose languages and easier to learn because the semantic gap between the user's mental model and the real model is smaller. Thus, from high-level, possibly visual models, a number of platform artifacts are generated which sometimes account for the 100% of the final application. Since the code generation process is automated from the models, there is a strong need to validate and verify them. In this respect, techniques for simulating the models (i.e. performing an animation showing their operational semantics) as well as to transform them into formal semantic domains for further analysis are of great interest.

Our work applies MDS techniques to the DL domain. In [11] we presented a DSL called VisMODLE oriented to the description of DLs, and built a modelling tool for it with an integrated code generator. VisMODLE is visual, which makes easier for people to learn and interpret; it is domain-specific, leaving less room for misunderstanding; and is formal, making possible the simulation and validation of the models. It allows the specification of the different aspects of a DL using four dimensions (or diagrams): services, collections, structure and society (i.e. interactions of services and actors).

In this work, VisMODLE is extended with a new type of diagram to specify behaviour (based on state machines). The operational semantics (i.e. simulator) of the new diagram type is formalized by using graph transformation [5]. This simulator allows a visual animation of the models (i.e. seeing messages being produced/consumed by actors and services) in order to validate the DL design and understand its behaviour. We have also designed a transformation from VisMODLE to Petri nets [13], a formal notation useful to specify concurrent and distributed systems. Its analysis techniques allow investigating system properties such as deadlocks, state reachability and invariants.

Paper organization. Section 2 introduces VisMODLE. Section 3 shows our MDS approach for generating DLs. Sections 4 and 5 describe the simulation and analysis techniques we propose, and its application to VisMODLE. Section 6 presents related work, and section 7 ends with the conclusions and future work. As a running example, we develop a simple DL for a university library.

2 VisMODLE, a DSL for Digital Libraries

VisMODLE is a DSL for the Visual Modelling of Digital Library Environments. A preliminary version of the language was presented at [11], but we have added a new diagram type for expressing behaviour. In this way, the specification of a DL in VisMODLE encompasses now five complementary dimensions: multimedia information supported by the DL (Collection Model); how that information is structured and organized (Structural Model); the services of the DL (Service Model); the societies of actors and services that interact to carry out the DL behaviour (Societal Model); and the individual behaviour of actors and services (Behavioural Model). Figure 1 shows the complete meta-model.

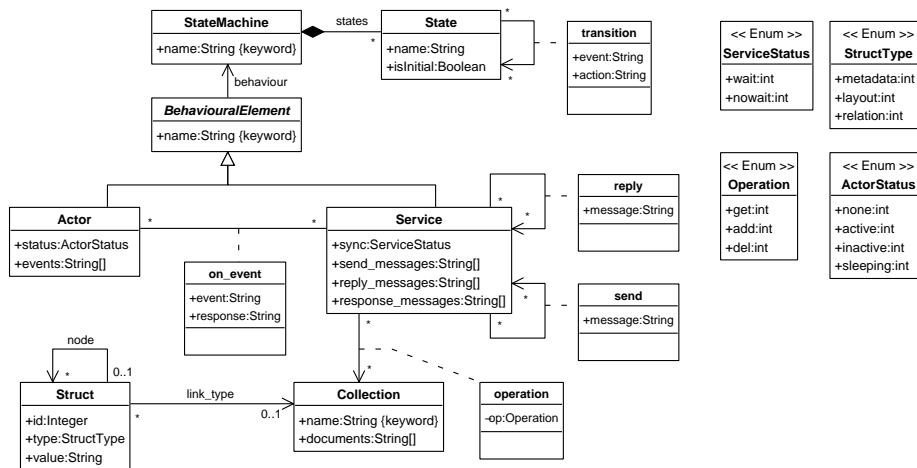


Fig. 1. The VisMODLE Meta-model.

Collections are sets of static (e.g. text) and dynamic (e.g. video) elements. Our running example includes a collection model (not shown for space constraints) with a collection (called *Library*) of two documents: long1.pdf and long2.pdf. The only relevant entity for this kind of diagram in the meta-model is class *Collection*.

The *Structural* diagram specifies how parts of a whole are arranged or organized. Structures can represent hypertexts, taxonomies, system connections, user relationships and containment. The window to the right in Figure 2 shows the structural model for the running example. Thus, collection *Library* is made of documents structured with *Publication*, *Author* and *Title* metadata information. The meta-model specifies that metadata entities (class *Struct*) can be connected together with the *node* relation (organized as a tree) and linked to a collection by a *link_type* relation.

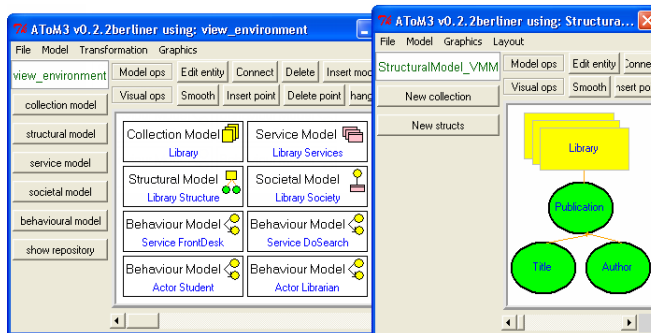


Fig. 2. Modelling Environment for VisMODLE.

Services describe activities, tasks, and operations (the functionality). Human information needs, and the process of satisfying them in the context of DLs, are well suited to description with services, including fact-finding, learning, gathering and exploring. In this diagram type, we describe the kind of messages services can produce and consume, and their synchronization. The only relevant entity from the meta-model for this diagram is class *Service*. In our example, we make two services available: *FrontDesk* and *DoSearch*. The former is responsible for managing communications between actors and is asynchronous (*sync* attribute is set to *nowait*), while the latter executes queries on the DL and is synchronous.

A *Society* is a set of entities and their relations. The entities include actors as well as hardware and software components, which either use or support services. Our meta-model is CSCW (Computer Supported Cooperative Work) oriented and thus includes the explicit description of the communication between entities and services. In complex systems such as DLs it is important to consider also the actors involved in the usage and so we introduced the society concept. A society is the highest-level view of a DL, which exists to serve the information needs of its entities and to describe the context of its use. *Society* diagrams are similar to UML collaboration diagrams. Figure 3 shows the society model for our example involving actors *Student* and *Librarian*. The scenario represents a *Student* borrowing a paper from the *Library*; he interacts with the *FrontDesk* service requesting the paper and obtaining a response message about its availability. The *FrontDesk* service forwards the borrow request to the *Librarian* actor. Then it sends a doc request message to the *DoSearch* service, which queries the document collection (*get* operation) using metadata information provided by the borrow request, and waits the result to send back the response. The service returns an *is_available* boolean message which is propagated as a response to the *Librarian* and eventually to the *Student*.

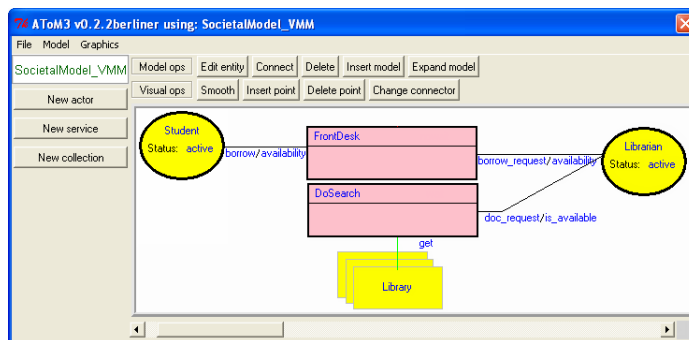


Fig. 3. A *Societal* Model.

Behavioural diagrams are used to specify the individual high-level activity of both services and actors by means of state machines. The transitions of the state machines are fired whenever a certain event occurs, which corresponds to the arrival of a message to the actor or service. In the transition, an action can also be specified, which is the sending of another message. Figure 4 shows the behavioural diagram for the

Librarian (initially in state *active*). If he receives an event *borrow_request*, he sends a *doc_request* message to service *DoSearch* and changes his state to *borrow_request*. Then, if he receives an event *is_available*, he becomes *active* again and sends the *availability* to service *FrontDesk*. For the current DL example, we have defined four different behavioural models, for actors *Student* and *Librarian* as well as for services *FrontDesk* and *DoSearch*.

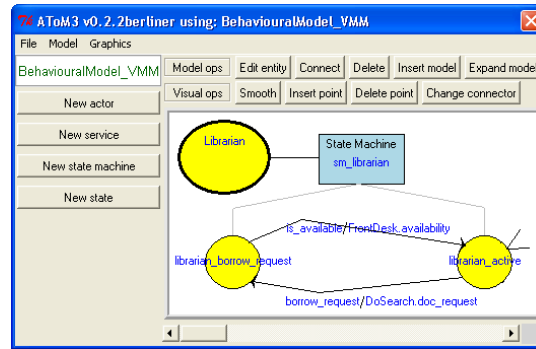


Fig. 4. A Behavioural Model.

3 Model Driven Approach to Digital Libraries

The overall architecture of our MDS approach for DLs is shown in Figure 5. The upper part depicts the process of designing the DSVL VisMODLE with the help of experts in the field of DLs, and its implementation in the meta-modelling tool ATOM³.

The environment for VisMODLE was defined by using the meta-modelling capabilities of ATOM³. Figure 6 shows a step in its definition. The window at the background (labelled “1”) partially shows the complete VisMODLE meta-model. The tool allows splitting the meta-model into different diagram types (called *viewpoints*). The five VisMODLE viewpoints are shown in the window labelled “2” in the figure. Window “3” contains the portion of the meta-model belonging to the *Structural* viewpoint. A special viewpoint named *repository_DigitalLibrary* contains the complete meta-model. In this way, in the generated environment, the user builds instances of the different viewpoints, and a repository model is created in the background with the gluing of the different diagrams the user has built. Consistency relations (shown as arrows in window “2”) specify how the different diagram elements are copied into the repository, how the different diagrams are kept consistent and how changes are propagated to the other diagrams if necessary. These arrows contain triple graph grammar (TGG) rules, which are automatically generated by ATOM³ [8]. In addition, a semantic view has been defined that allows expressing the semantics of the repository by using a transformation into Petri nets (see section 5). This is also performed by using TGG rules.

With the information described before, ATOM³ generated a customized modelling environment for VisMODLE (shown at the background of Figure 2). DL designers can

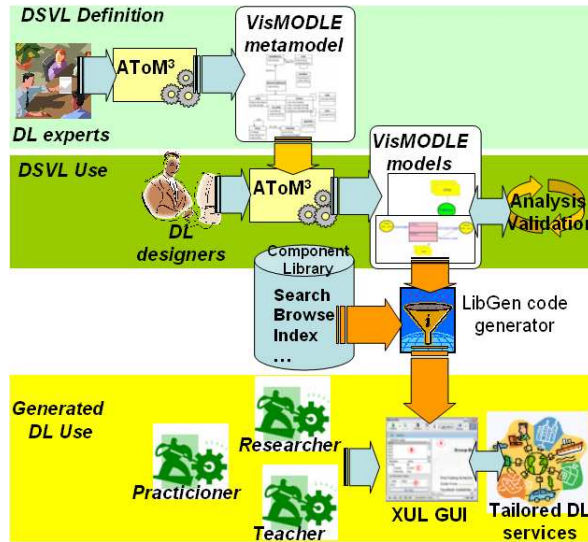


Fig. 5. The MDS Architecture for DLs.

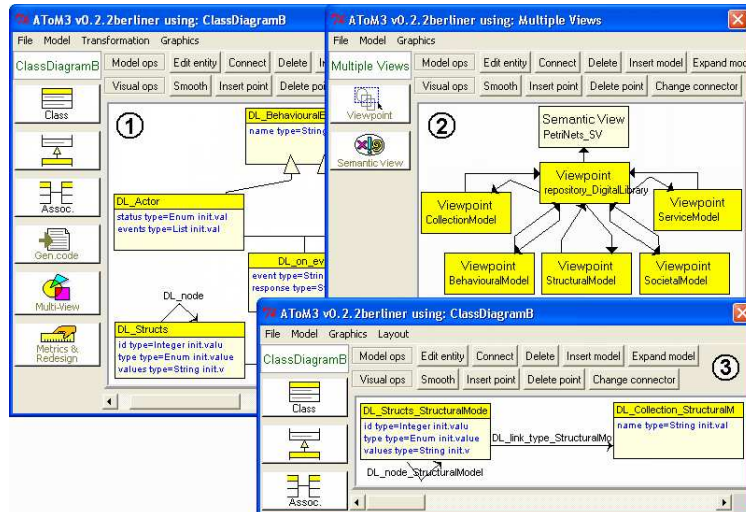


Fig. 6. Building the Environment for VisMODLE.

use this environment to build VisMODLE models (layer “DSVL use” in Figure 5). The environment also integrates a simulator that allows the visual animation of the models. The simulator helps to validate and understand the DL design by observing the message flow between services and actors (see section 4). As stated before, some analysis capabilities have been also integrated based on model transformation into Petri nets (see section 5).

Finally, we have provided the environment with a code generator (called “LibGen” in Figure 5) able to produce XUL code for the user interface (UI) of the DL. The generator also selects predefined service components to implement the required functionalities of the VisMODLE models [11]. Figure 7 shows the generated interface for the example (for the *Librarian*). The generated UI is built upon a set of XUL template files that are automatically specialized depending on the attributes and relationships designed in the modelling phase. The layout template for the UI is divided into two columns. The left part manages the collections of documents and its metadata information. The right part manages visualization and multimedia information obtained from documents. The basic features provided with the UI templates are document loading and visualization, metadata organization and management.

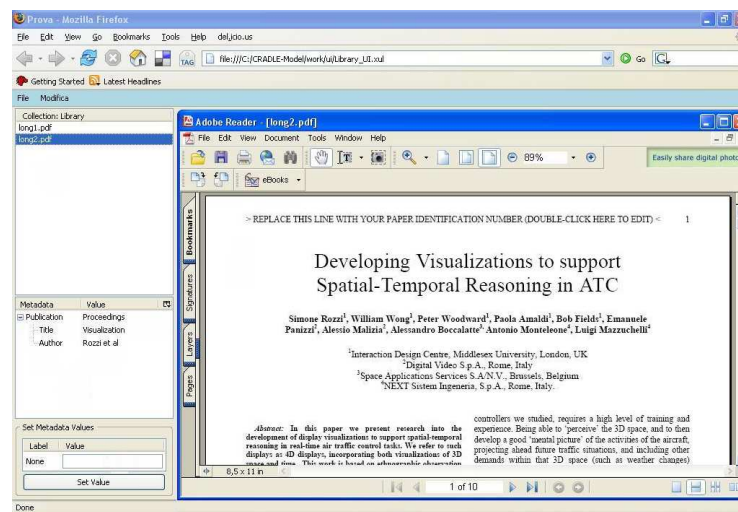


Fig. 7. Generated GUI for the Example DL.

For the composition of the services specified in the VisMODLE models, we have used *XDoclet* [17]. This is an open-source code generation library which enables Attribute-Oriented Programming for Java via insertion of special tags. It includes a library of predefined tags, which simplify coding for various technologies such as web services.

4 Simulation of DL models

In order to visualize and better understand the behaviour of a DL, we have built a simulator for VisMODLE. The simulation is performed in the repository, since it needs all the dynamic information expressed in VisMODLE (i.e. the societal and behavioural models). The simulation consists of the animation of the state machines that describe the behaviour of actors and services. In the simulation, it is checked that the events

posed of rules, with graphs in their left and right hand sides (LHS and RHS respectively). In order to apply a rule to a graph (called *host graph*), first a matching (an occurrence) of the LHS has to be found on it. Then, the rule is applied by substituting the match in the host graph by the rule's RHS. In addition, rules can define application conditions that restrict their applicability. One of the most used are the so-called Negative Application Conditions (NACs). These are graphs that must not be present in the host graph for the rule to be applied. On the contrary, Positive Application Conditions (PACs) are graphs that must be present in the host graph in order to apply the rule. Finally, we can combine meta-modelling and graph transformation allowing abstract nodes to appear in rules [5]. In this way, nodes can be matched to instances of any subclass, greatly improving the expressive power of rules.

Our simulator is made of five rules, three of them shown in Figures 9, 10 and 11. Rule in Figure 9 assigns an *input* element to each behavioural element, that is, to each actor and service in the repository (i.e. each possible concrete subclass of class *BehaviouralElement*). Initially, the input does not contain any message, and points to the initial state of the behavioural element's state machine. If the behavioural element already has an associated input element (NAC), the rule is not applied. In this way, we are sure there is only one input for service and actor.




Fig. 9. Simulation Rule *create input*.

The other four simulation rules perform a simulation step (i.e. a state transition if the required event was produced) for different cases. For example, the rule in Figure 10 considers transitions where no action is defined. In this case, if a behavioural element receives a message whose label matches some of its outgoing transitions (LHS), then the message is processed and the current state is changed (RHS). The attribute condition in the LHS checks that the message name in the input is equal to the event specified in the outgoing transition of the current state.

Similarly, rule in Figure 11 considers transitions with an action that sends a message to another behavioural element (i.e. the action has the form “target_element.message”). Here, in addition to the previous conditions for a behavioural element to change its state, it is necessary that the target element accepts messages of that kind. This is given by the PAC below, which checks that the corresponding link has been specified in the societal model. Indeed, one of these three positive conditions has to be fulfilled in order to be able to apply the rule: either the source element is an actor that sends a message to a service (first case), or it is a service that sends a message either to an actor (second case) or to another service (third case). The use of abstract nodes (labels 8 and 11) together




Fig. 10. Simulation Rule *process event*.

with the three PACs allows using a single rule for the three cases. This rule deletes the message from the input, changes the current state and creates a message in the target element input. The action below the RHS assigns the message specified in the action's transition as the name of the newly created message.

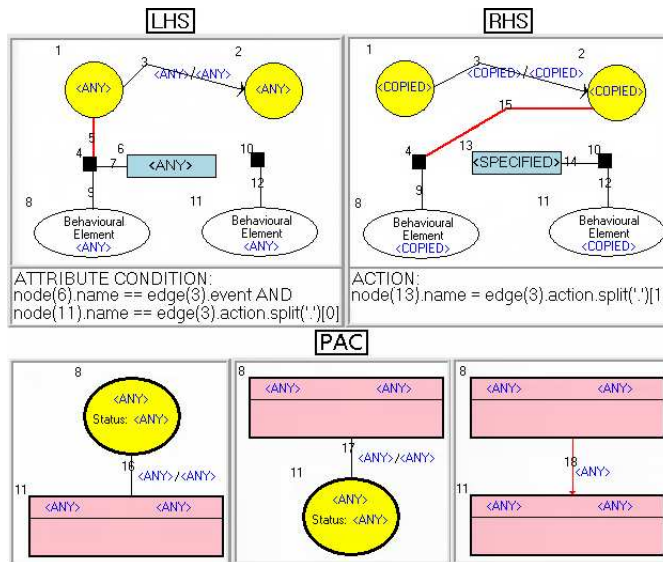


Fig. 11. Simulation Rule *process event with action*.

Finally, two similar rules (not shown in the paper) consider transitions where either an action is performed without the necessity of an event, or a transition occurs with neither event nor action associated.

5 Analysis of DL models

We have provided the VisMODLE environment with some analysis mechanisms to validate the dynamics of a DL specification. More in detail, given the set of models that conform a DL design, we can:

1. check if some behavioural element reaches a deadlock state (i.e. local deadlock). In the case of services, we normally want to keep them always available, so they should not define final states. Final states are shown highlighted in the model as result of the analysis, and they are presented textually in a dialog window as well.
2. check whether a given behavioural element reaches certain state. The sequence of state transitions leading to the requested state is shown highlighted as a result, and also textually.
3. ask if a given message is always sent in any possible execution flow. The answer (*true* or *false*) is shown in a dialog window.
4. check if the system execution always finishes (i.e. global deadlock).
5. detect if there is some state for a behavioural element that is never reached in the given society. This may be considered a design error, since we should not define states that are not possible.
6. check if some behavioural element can receive an unbounded number of messages at some point, which could lead to an overflow.

In order to provide these analysis mechanisms, we have expressed the operational semantics of the models by means of Petri nets [13] (using a semantic view, see Figure 6). Thus, we have defined a triple graph transformation system (TGTS) [7] that transforms the repository into the equivalent Place/Transition Petri net. Once the net is obtained, we internally use analysis techniques based on the reachability/coverability graph, as well as model-checking [3]. Similarly to graph grammars, TGTSs are made of rules that contain, in this case, triple graphs (instead of simple ones). Triple graphs are made of three different graphs: source, target and correspondence. In our case, the source graph is the repository, the target graph is the Petri net resulting from the transformation, while the correspondence graph contains elements relating the elements in the other two graphs.

Figures 12 and 13 show some rules of the TGTS that build the Place/Transition Petri net from the VisMODLE repository.

The idea of the transformation is transforming the state machines of each behavioural element into what we call Petri net modules. With this purpose, states are translated into places (triple rule shown in Figure 12) and transitions between states are translated into transitions between places (triple rule in Figure 13). If a state is initial, then the associated place will contain one token. In other case, the place will be empty. In addition, an extra place is created for each possible message invocation (i.e. for each possible event and action specified in each state machine). These places are the *interface* of the Petri net module of the behavioural element. If a transition needs an event for being executed, then an arc is created from the place corresponding to the event message, to the Petri net transition corresponding to the state machine transition. In that case, the Petri net transition can be executed only if an event of that type is received




Fig. 12. Triple Rule SM state to PN place.




Fig. 13. Triple Rule SM transition to PN transition.

(i.e. a token is in the right interface place), and in addition, the behavioural element is in the right state (i.e. a token exists in the place corresponding to the state source of the transition). Similarly, if a transition specifies an action, an arc is created from the Petri net transition to the right place of the behavioural element's interface specified in the action. In this way, executing the transition implies creating a token in the right interface place, and results in the interconnection of the different modules. Figure 14 shows the resulting Petri net after applying the TGTS to the DL example.

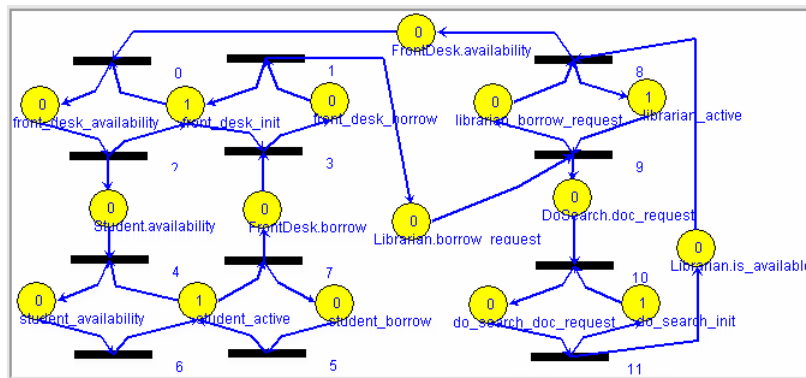


Fig. 14. Place/Transition Petri Net for the DL Example.

Each analysis provided to the VisMODLE modelling environment implies evaluating a CTL logical expression on the net's coverability graph [3]. With this purpose, we use a function that calculates the coverability graph of the net, as well as a model checker to evaluate the expression, both implemented in AToM³. In order to use these analysis techniques, users of the VisMODLE environment do not have to know Petri nets at all. All the analysis process is hidden. CTL expressions are defined when the modelling environment is generated, and hidden to the final user. The results of analysing the Petri net are back-annotated and shown to the user in the VisMODLE notation, which is the notation that he knows. This is possible since we maintain in the correspondence graph the relations between the elements in the DL and the ones in the Petri net. AToM³ provides declarative back-annotation mechanisms that allow the specification of the elements to highlight in the source model as a result of the analysis.

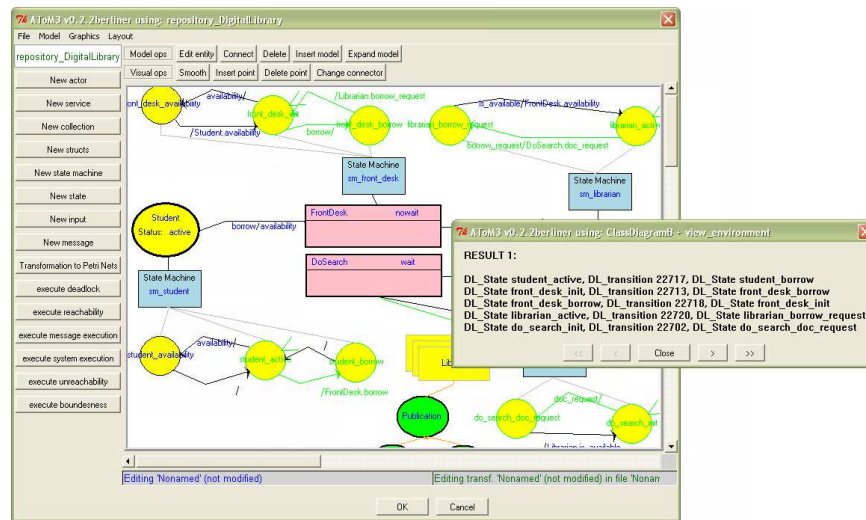


Fig. 15. Execution of Analysis Method Call.

For example, Figure 15 shows the repository interface of the generated VisMODLE environment. The repository interface includes one button for each one of the defined analysis method calls. In order to check a property, the end-user has to open the repository and click on one of these buttons. Internally, the tool performs the transformation to the semantic domain, evaluates the CTL formula that corresponds to the property on the Petri net's coverability graph, and returns the result according to the defined back-annotation mechanism. In the figure, the analysis called *reachability* for state *front_desk_borrow* is shown. The window to the right allows navigating through the different solutions (i.e. the different path executions) that lead to such state. The VisMODLE elements that conform a single solution are shown in the navigator and highlighted in the model. In this way, the analysis result is shown in terms of the original notation to the user, who does not need to have knowledge about any formal method. For the example DL, the net showed neither global nor local deadlocks, all the states

were reachable, and the tool signalled the possibility of unbounded number of *borrow* and *borrow_request* messages.

6 Related Work

Formal models for DLs are rarely found, likely due to the complexity of the field. The approach of [16] defines DLs as a combination of a special-purpose database and a hypermedia-based UI, and formalizes them by using the Z language. In [2] a multidimensional query language for DLs is used that rely on the notions of views and versions, metadata formats and specifications, and a first-order logic based language. [6] presents a formal foundation theory on DLs, called 5S, based on streams, data structures, spaces, scenarios and societies. We were inspired on it in order to design VisMODLE. However, these formal approaches do not formally specify how to derive the DL implementation from the DL model.

Other declarative approaches, such as the Digital Library Definition Language [12], the METIS framework [1] and the FEDORAs structoid approach [4], are not supported by a strict underlying formal theory.

To the best of our knowledge, none of these approaches provide a customized environment supporting code generation as well as validation (animation) and analysis techniques. By using MDS techniques, we help DL experts to cope with the complexity of DL designs without dealing with coding. In addition, the use of a DSVL makes it easier for people to learn the domain concepts and interpret the models, and allows modelling interactions among DL systems and users (as proposed in the HCI field). Finally, our approach generalizes some metadata schemas such as DC, in the sense that DC data structures can be modelled by using our structure entities and their relationships.

7 Conclusions

In this paper, we have presented VisMODLE, a DSVL for building DLs in a model-driven way. The language is made of five diagram types (services, collections, structure, society and behaviour) to describe the aspects of a DL. We have generated a modelling environment for it with the AToM³ tool. The environment integrates a code generator in order to produce the DL. Although it is in its alpha version, we have already used it to build prototypes. In addition, we have built a visual simulator using graph transformation to graphically validate the DL behaviour at the model level. In order to perform further analysis, a transformation into Petri nets has been designed, which allows checking model properties, such as reachability or deadlocks. This analysis is made by internally performing model-checking of the coverability graph by using predefined temporal logic formulae.

In the future, we intend to take into account scalability issues, as well as to generate code for some orchestration language to reflect the overall behaviour of the societal model.

Acknowledgements. Work sponsored by projects MODUWEB (TIN2006-09678) and MOSAIC (TIC2005-08225-C07-06) of the Spanish Ministry of Science and Education.

References

1. Anderson, K. M., Andersen, A., Wadhvani, N., Bartolo, L. M.: Metis: Lightweight, flexible, and web-based workflow services for digital libraries. *JCDL*, 98-109 (2003)
2. Castelli, D., Meghini, C., Pagano, P.: Foundations of a multidimensional query language for digital libraries. *ECDL*, 251-265 (2002)
3. de Lara, J., Guerra, E., Vangheluwe, H.: Metamodelling, graph transformation and model checking for the analysis of hybrid systems. In: *AGTIVE03*. LNCS, vol. 3062, pp. 292-298. Springer (2003)
4. Dushay, N.: Using structural metadata to localize experience of digital content. *CoRR*, cs/0112017 (2001)
5. Ehrig, H., Ehrig, K., Prange, U., Taentzer, G.: *Fundamentals of Algebraic Graph Transformation*. Springer-Verlag (2006)
6. Gonçalves, M. A., Fox, E. A., Watson, L. T., Kipp, N. A.: Streams, structures, spaces, scenarios, societies (5s): A formal model for digital libraries. *ACM Trans. Inf. Syst.* 22(2), 270-312 (2004)
7. Guerra, E., de Lara, J.: Attributed typed triple graph transformation with inheritance in the double pushout approach. In: *Technical Report UC3M-TR-CS-06-01*, Universidad Carlos III de Madrid (2006)
8. Guerra, E., de Lara, J.: Model view management with triple graph transformation systems. In: *ICGT06*. LNCS, vol. 4178, pp. 351-366 (2006)
9. Kochtanek, T. R., Hein, K. K.: Delphi study of digital libraries. *Inf. Proc. Manag.* 35(3), 245-254 (1999)
10. Licklider, J. C. R.: *Libraries of the Future*. MIT Press, Cambridge, Mass (1965)
11. Malizia, A., Guerra, E., de Lara, J.: Model-driven development of digital libraries: Generating the user interface. In: *Proc. MDDAUI06* (2006)
12. Maly, K., Zubair, M., Anan, H., Tan, D., Zhang, Y.: Scalable digital libraries based on nestr/dienst. *ECDL*, 168-179 (2000)
13. Murata, T.: Petri nets: Properties, analysis and applications. In: *Proc. of the IEEE* 77(4), 541-580 (1989)
14. Reddy, R., Wladawsky-Berger, I.: Digital libraries: Universal access to human knowledge report to the president. In: *PITAC, Panel on DLs* (2001)
15. Völter, M., Stahl, T.: *Model-Driven Software Development*. Wiley (2006)
16. Wang, B.: A hybrid system approach for supporting digital libraries. *JDL* 2(2-3), 91-110 (1999)
17. XDoclet home page: <http://xdoclet.sourceforge.net/xdoclet/index.html>