Automated Analysis and Code Generation for Domain-Specific Models

George Edwards
Blue Cell Software
george@bluecellsoftware.com

Yuriy Brun
University of Washington
brun@cs.washington.edu

Nenad Medvidovic
University of Southern California
neno@usc.edu

Under Review

***Abstract*—Domain-specific languages (DSLs) concisely express the essential features of system designs. However, using a DSL for automated analysis and code generation requires developing specialized tools. We describe how to create model analysis and code generation tools that can be applied to a large family of DSLs, and show how we created the LIGHT platform, a suite of such tools for the family of software architecture-based DSLs. These tools can be easily reused off-the-shelf with new DSLs, freeing engineers from having to custom-develop them. The key innovation underlying our strategy is to enhance DSL metamodels with additional semantics, and then automatically synthesize configurations and plug-ins for flexible analysis and code generation frameworks. Our evaluation shows that, for a DSL of typical size, using our strategy relieves software engineers of developing approximately 17,500 lines of code, which amounts to several person-months of programming work.**

I. Introduction

Some software-intensive systems, such as aerospace systems and sensor network applications, require rigorous design modeling. *Domain-specific* modeling tools and languages allow meticulous design analysis and automatic code generation, improving system quality. In contrast to standardized modeling languages like UML, domain-specific languages (DSLs) allow engineers to focus on the design decisions relevant to the domain and use the most suitable concepts and abstractions.

Today's model-driven engineering (MDE) platforms, such as the Generic Modeling Environment (GME) [15] and the Eclipse Graphical Modeling Framework (GMF) [16], ease the creation of custom model editors for DSLs. Software engineers only need to define a *metamodel* — a formal specification of a DSL — and these platforms automatically synthesize a model editor that uses the DSL's symbols and enforces its syntax.

Nevertheless, industry adoption of domain-specific modeling technologies has been more limited than that of standardized modeling solutions, particularly UML. One key reason for this disparity is that DSL *analysis* and *code generation* tools (often referred to as *model interpreters*) must be constructed manually. Meanwhile, UML-based analysis and code generation tools are available off-the-shelf. While domain-specific tools can perform more targeted analysis and more complete code generation, the difficulty of tool creation and maintenance reduces the appeal of domain-specific modeling, particularly for small- and medium-scale software systems [26].

In this paper, we present a solution that greatly reduces the costs of model interpreter creation and maintenance. We show how specific enhancements to metamodels can allow an MDE platform to automatically synthesize analysis, simulation, and code generation tools, just as existing MDE platforms automatically synthesize model editors. While our approach applies to DSLs in general, in this paper we focus on software architecture-based modeling. The field of software architecture is characterized by a large number of DSLs, arising from different architectural styles, design patterns, modeling notations, analysis tools, middleware platforms, and frameworks [35]. We implement our approach in an MDE platform called LIGHT (Leveraging Isomorphism to Generate Heterogeneous DSL Toolchains). LIGHT allows engineers to customize architecture-based DSLs for their particular project modeling needs and then automatically synthesize (1) a model editor, (2) a system simulator for analyzing latency, memory usage, energy consumption, and reliability, and (3) a middleware-based system implementation. LIGHT eliminates a substantial amount of tool building and maintenance work required by existing MDE platforms.

Consider a team designing the software module for a moon-landing spacecraft. The team wants to (1) create models to formalize their designs, (2) evaluate those designs in simulation, and (3) automatically generate code from the model to ensure the implementation conforms with the model. Recognizing the highly specialized nature of the software and the need to focus on design decisions specific to the domain, the team decides to use an architecture DSL.

Using existing MDE tools, the team specifies the DSL metamodel and uses the generated editor to develop a series of architecture models. They then develop a custom model interpreter that translates their domain-specific models into simulation code. This interpreter implements a complex transformation function, akin to a compiler, which realizes the semantics of the DSL. Every time the engineers alter the DSL, they must update the model interpreter. As we discuss in Section V, this coding effort can amount to as much as four person-months of work for a DSL of moderate size and complexity. The team must exert the same significant effort to create the other desired tools. This cost reduces the appeal of DSLs and drives engineers to use standardized languages, even if they are sub-

optimal for the modeling task at hand.

In contrast, using LIGHT, the team only needs to specify a slightly expanded DSL metamodel, and LIGHT generates the model editing, simulation, and code generation tools automatically from the metamodel. If the team makes changes to the DSL, the tools are automatically updated to conform to the new semantics. As we discuss in Section V, the effort required to create the expanded metamodel is minor compared to the effort of creating model interpreters.

The insight that allows LIGHT to automate model interpreter creation and maintenance is that model editors and model interpreters are *isomorphic* [10]: rendering models within an editor is just another form of model interpretation. The implication of this insight is that model editors and interpreters can be treated as analogs conceptually and architecturally — a hypothesis we set out to validate by designing, implementing, and evaluating LIGHT.

We previously studied modularizing domain-independent and domain-specific model interpretation logic, an approach we implemented in the XTEAM platform [7]–[9]. XTEAM promoted reuse of significant portions of model interpreter implementations, but still required engineers to manually program certain domain-specific logic. In contrast, LIGHT leverages the isomorphism and generates model interpreters using the same method that has already proven so successful for the construction of model editors.

The contributions of our work are:

- a novel architecture for creating modeling, analysis, and code generation toolchains from DSL metamodels,
- LIGHT, an instantiation of the architecture targeted at architecture-based software development,
- an evaluation of LIGHT's utility across nine different architectural DSLs, and
- an elaboration of the trade-offs associated with LIGHT.

The remainder of this paper is organized as follows. Section II outlines the problem we aim to solve. Section III describes our solution and its applicability. Section IV details the implementation of the LIGHT platform. Section V evaluates LIGHT. Section VI compares our approach to related work. Section VII concludes the paper.

II. The DSL Model Interpreter Problem

Throughout this paper, we will rely on Lunar Lander (LL), the moon-landing spacecraft software introduced in Section I. LL has been used as an instructional tool for software architecture concepts, via its many variations [35]. Suppose an engineering team is tasked with building LL. They analyze the system requirements and arrive at a set of high-level design goals, including: components should be independent (facilitating reuse) and dynamically configurable (enabling runtime system adaptations), and critical components should be replicable (trading off efficiency for reliability).

Based on these goals, the team decides to employ the Myx [5] architectural style. Myx supports layered architectural composition and is aimed at flexible construction of distributed systems, ensuring component decoupling and enabling dynamic adaptation. Furthermore, the team elects to rigorously model candidate designs to (1) document the system architecture, (2) evaluate alternative designs with respect to latency and reliability, and (3) generate code for the system implementation. For example, the team would like to use LL models to explore alternative component replication strategies in an off-the-shelf simulation tool.

Recognizing LL's highly specialized nature and the need to focus on design decisions specific to the domain, the team decides to use a DSL to model candidate designs. Using a DSL frees the team from the constraints and feature bloat of a standardized language like UML, allows them to customize the look-and-feel of their diagrams, and allows a single model to contain all the relevant information, such as the parameters necessary for latency and reliability analysis.

The most straightforward approach is for the team to develop their own model editing, analysis, and code generation tools that are customized for the Myx-based DSL and the particular semantics the team desires. For all but the simplest DSLs, developing these tools would require writing tens-to-hundreds of thousands of lines of code.

A much better approach is to use an MDE platform, such as GME or GMF. To do so, the team first specifies a metamodel for their DSL in a provided metamodel editor. The metamodel encodes the rules of the Myx style and defines the necessary latency and reliability analysis parameters. The metamodel conforms to the MDE platform's metamodeling language (or *metalanguage*) and consists of instances of the metalanguage types (called *metatypes*). Each metatype instance captures the type definition for a domain-specific modeling element. For example, the team can define a `MyxLink` type in their DSL that represents an association between component interfaces. In a GME metamodel, the `MyxLink` type could be an instance of the `Connection` metatype. In the team's models of candidate LL designs, individual associations between components are then represented by instances of the `MyxLink` type.

Using the metamodel, existing MDE platforms automatically create a custom model editor by generating configuration files or plug-ins for a *model editor framework*. The editor renders the model and allows model manipulations while enforcing the DSL constraints. For example, the instances of the `MyxLink` type can be rendered as dashed lines. Defining a metamodel using an MDE platform in this way allows the LL development team to automatically generate a custom graphical model editor for their DSL.

However, existing MDE platforms provide only a partial remedy: they do not provide built-in analysis and code generation support for DSLs. Instead, they require software engineers to implement custom model interpreters. For example, if the LL development team wants to use an off-the-shelf simulation tool (e.g., MATLAB) to analyze their candidate designs, they must implement a model interpreter that translates their

Under Review

DSL models into the form required by the tool. Similarly, to generate executable code for a runtime platform (e.g., .NET), the team must implement a model interpreter to produce that code. MDE platforms cannot automate this process because the metatypes provided by these platforms lack sufficient semantics to be automatically mapped to other forms. Their semantics are limited to those needed to synthesize model editors (i.e., map metatypes to graphical display elements).

As we show in Section V, creating the interpreter to generate executable simulations from Myx models requires the LL team to write over 17.5K non-trivial lines of code. Moreover, after selecting a preferred design, the team may want to generate source code to ensure that the implementation conforms to the model. In that case, they have to write thousands of lines of additional, again non-trivial code to implement another interpreter that generates the implementation for the desired middleware platform or OS.

Thus, the consequences of existing MDE's limitations are:

- Analysis and generation tools must be manually constructed, which is difficult and expensive [19], [31].
- Language and tool reuse is hard, as using them in new engineering contexts can require significant rework.
- The maintenance of domain-specific analysis tools and code generators is burdensome, as they must be updated whenever the corresponding DSL changes.

III. Automated Synthesis of Model Interpreters

In this section, we describe our approach to building an MDE platform. The approach is aimed at drastically reducing the time and effort required to create end-to-end domain-specific modeling, analysis, and code generation toolchains. Our guiding idea is to leverage additional semantics within metamodels to enable automatic generation of configuration files and plug-ins for extensible analysis and code generation frameworks. Our approach comprises three activities: (1) the MDE platform developers select the analysis and code generation capabilities they want supported by the platform, (2) they extend the semantics of the platform metatypes to enable those capabilities, and (3) they create a *metainterpreter* and *model interpreter framework* that together perform the automated synthesis of tools that implement those capabilities. Note that the process described only needs to be carried out by MDE platform *developers*; MDE platform *users* (such as the LL development team) can then obtain the MDE platform off-the-shelf and automatically generate tools for their DSL.

A. Capability Selection

The developers of an MDE platform must first select which analysis and code generation capabilities to support, based on the expected usage of the platform. This decision is fundamental to our proposed MDE platform architecture because it determines which semantics the metamodels will include. We do not address the process of making this decision in this paper and instead focus on how to implement an already chosen set of capabilities. For our LIGHT reference implementation of the approach, we chose to support a model editor, a middleware platform, and a simulation engine that analyzes latency, memory usage, energy consumption, and reliability.

B. Metatype Semantics Extension

Whereas today's approaches require engineers using a DSL to manually build model interpreters that encode semantics, our approach embeds these semantics in the DSL metamodel. We describe these metamodels in Section IV-A. A set of metainterpreters then automatically generates the model interpreters from the metamodel. We describe the metainterpreters in Section IV-B. Since the metamodel is crucial for the interpreter generation, it is important to determine the exact semantics the metamodel needs to capture. In order to automatically synthesize model interpreters, the metamodel must define a complete *semantic mapping* from domain-specific elements to target platform elements (e.g., graphical display elements in the case of an editor or programming language constructs in the case of a code generator).

For example, consider the `MyxLink` DSL type from the Myx metamodel developed by the LL team. In an existing MDE platform such as GME, this type might be defined by an instance of the `Connection` metatype. GME is pre-programmed with the *presentation semantics* (model editor behavior) of instances of the `Connection` metatype: drawing a line between two objects. Thus, GME maps `MyxLink` instances in LL models to classes that render a dashed line. GME does not know, e.g., the *simulation semantics* of instances of the `Connection` metatype. On the other hand, using our approach, as implemented in LIGHT, `MyxLink` is an instance of a specially defined `Link` metatype that includes presentation semantics as well as simulation and *code generation semantics*.

Our approach directly attaches semantics to metatypes as (1) semantic *assumptions*, which are behavior definitions that hold for all domain-specific types that are instances of that metatype, and (2) semantic *properties*, which are typed attributes and associations with other metatypes that allow an engineer to select among various behavior options. For example, in LIGHT, a semantic assumption is that all domain-specific types defined by instances of the `Link` metatype exhibit the behavior of transferring data between two other entities. A semantic property of the `Link` metatype, called `capacity`, allows engineers to specify, in the metamodel, whether the data transfer channel can become full and the resulting behavior.

This approach has the advantage that metamodel developers do not need to write intricate formal semantic specifications. However, the metatype assumptions and properties are chosen by the metalanguage designers, thus the space of possible semantics that can be captured is fixed. This results in a trade-off between the ability to synthesize supporting toolsets and DSL flexibility. Attaching additional semantics to metatypes increases the space of possible model editors, analysis engines, and code generators that can be synthesized. For example,

Under Review

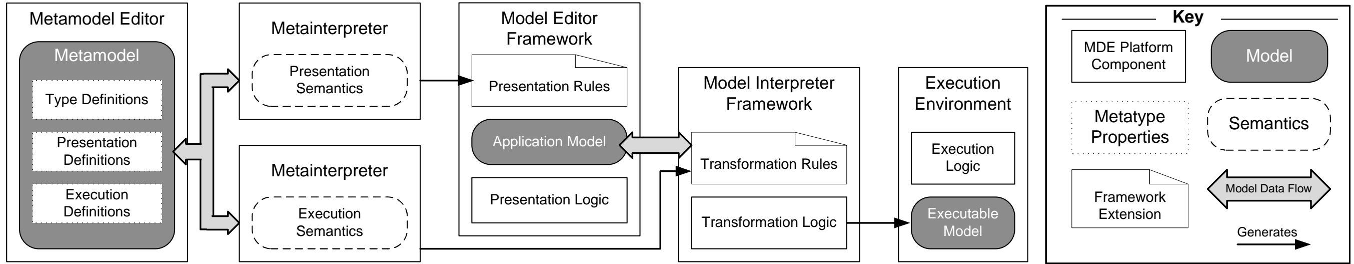

Figure 1. Our approach to automated tool generation for DSLs.

current MDE platforms only synthesize model editors because their metatype semantics include only visualization and editing concerns. On the other hand, attaching additional semantics to metatypes decreases the space of possible DSLs that can be specified in a domain-specific modeling platform. For example, if metatypes include fixed semantics for graphical rendering and editing, they cannot be (easily) used to specify a textual language. Section IV-A discusses the semantics captured in LIGHT's metalanguage.

C. Interpretation Component Development

Our approach fundamentally differs from existing approaches, including our own previous work, in its mechanism for model interpretation. The approach does not require any manually coded components to perform system analysis or to generate executable code. Figure 1 depicts the roles and interactions of the MDE platform components that participate in analysis and code generation. Each interpretation capability is implemented through a paired *metainterpreter* and *model interpreter framework* (MIF). As noted earlier, we have implemented three such interpretation capabilities in LIGHT (a model editor generator, a simulation generator, and a code generator) using this general architecture. We summarize the function of metainterpreters and MIFs here, while Section IV-B discusses their details.

Existing MDE platforms lack MIFs and their associated metainterpreters. Metainterpreters (1) take as input a metamodel, including the metatype properties needed to map domain-specific models to a target platform, (2) use those metatype properties to derive domain-specific type semantics, and (3) determine a set of rules for transforming each domain-specific type to the analysis, simulation, or execution platform language. Metainterpreters encode the transformation rules in an automatically-generated *MIF extension* for a specific MIF; the MIF implements the actual transformation logic.

Each MIF is a template for a family of model interpreters. To be reusable, a MIF must encapsulate transformation logic or algorithms that (1) are useful in a wide variety of contexts, and (2) can be flexibly applied in different ways to achieve different semantics. A MIF is analogous to a virtual machine in that it provides and executes an instruction set composed of model transformation operations. Transformation steps that only depend on semantic assumptions (i.e., not metatype properties) are hard-coded into the MIF, while transformation steps that vary based on metatype properties are programmable via extension points.

A domain-specific type's semantics are a subset of all possible semantics permitted by the MDE platform; each possible semantic option corresponds to a different usage of MIF extension points. Therefore, a metainterpreter includes (1) a mapping of the metatype instance property values to a set of semantic definitions, and (2) a mapping of semantic definitions to a set of MIF extension point usages. A generated MIF extension, which may take the form of configuration files or plug-ins, modifies, extends, and controls the functionality of the MIF using extension points built into the MIF. The extended MIF converts a domain-specific model into an executable or analyzable program for the target platform.

IV. The LIGHT Platform

To verify the feasibility of our approach, we built the LIGHT MDE platform. LIGHT is intended for software architecture-based modeling, analysis, and code generation. LIGHT automatically generates model interpreters for any DSL defined by a LIGHT metamodel, by configuring a model interpreter framework (MIF) with a domain-specific MIF extension generated by a corresponding metainterpreter. Each such metainterpreter-MIF pair generates a different type of interpreter. The LIGHT platform contains three such pairs that target, respectively, GME's open-source model editor [15], a variant of the Adevs discrete event simulation engine [27], and the Prism-MW middleware platform [24]. This means that engineers using LIGHT can create domain-specific models in a customized model editor, and then analyze those models in the simulator and generate code for Prism-MW, all with virtually no tool-building overhead.

To use LIGHT, engineers first create a DSL metamodel via LIGHT's provided metamodel editor. Then, to generate the three interpreters, LIGHT first invokes the appropriate metainterpreter, which produces an MIF extension (a set of C++ plug-in classes) by deriving the simulation or implementation semantics of the DSL types in the metamodel. LIGHT then compiles the provided MIF (also implemented in C++) with the extension. The output of the compilation is a domain-specific model editor, simulation generator, or Prism-MW code generator, already configured with DSL's custom semantics.

Under Review

Next, we discuss LIGHT's metamodeling facilities (Section IV-A) and the implementation of LIGHT's model interpretation components (Section IV-B), and then focus on using simulations generated by LIGHT to analyze domain-specific models with respect to latency, reliability, and other system qualities (Section IV-C). The details of the Prism-MW code generator are elided for space.

A. *Metamodeling*

LIGHT maps DSL types to runtime objects (e.g., simulation and Prism-MW objects) through the use of a metalanguage enhanced with semantic assumptions and properties (recall Section III). We designed the LIGHT metalanguage in a two-step process. First, since LIGHT is intended for software architecture-based modeling, analysis, and code generation, we conducted a literature review to identify the common elements, abstractions, and patterns used for architectural models. We identified ten important metatypes: *architecture*, *component*, *resource*, *interface*, *link*, *implementation*, *operation*, *task*, *data type*, and *property*. Second, we defined the semantic assumptions and properties for each metatype. Recall that semantic assumptions are behavior definitions that are inherent to each LIGHT metatype and hold for all domain-specific types that are instances of that metatype. Semantic assumptions can be further classified as *capabilities* and *responsibilities*. Capabilities describe behaviors that instances of the metatype exhibit by default, while responsibilities describe behavior constraints that instances of the metatype must respect. Semantic properties are typed attributes and associations that capture semantic variations and options. Software engineers customize the semantics of a DSL by setting the values of semantic properties in a LIGHT metamodel.

Figure 2 depicts the ten LIGHT metatypes (and an `Entity` supertype) and their semantic properties. An `Architecture` defines a system as a collection of components, resources, and other types, and their relationships. A `Component` defines a set of interfaces and maps each interface to either the interface of a sub-component or an implementation. Each `Interface` defines a component interaction point. `Implementations` capture computational logic and state in terms of sequences of instructions (e.g., methods) or state-transition systems (e.g., Finite State Processes). `Links` represent logical connections among interfaces used to exchange information and control. `Resources` are entities provided by the execution environment that components use to perform tasks. `Operations` define component service access points in terms of data and control exchange and are grouped together within interfaces. `Tasks` are reusable units of functionality within a component. Finally, `DataTypes` represent objects exchanged between components or maintained as part of a component's state.

To understand the space of possible DSLs that can be captured in a LIGHT metamodel (and for which a custom model editor, simulation generator, and code generator can be automatically synthesized) is it necessary to understand the semantic assumptions and properties of the LIGHT metatypes. Any language whose semantics are (1) compatible with the LIGHT semantic assumptions, and (2) within the semantic options provided by the LIGHT metatype properties can be described. Due to space constraints, we cannot enumerate and explain the semantic assumptions and properties of all ten metatypes here. Instead, we will highlight three LIGHT metatypes: `Component`, `Interface`, and `Resource` (see Figure 3). We elaborate on these three types next; we refer the reader to [22] for a complete table of all metatypes and [12] for their detailed descriptions.

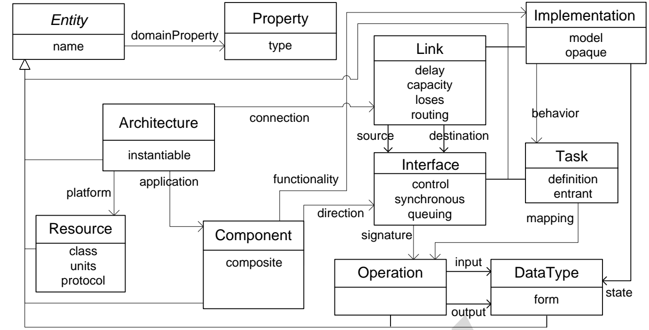

Figure 2. Summary of the LIGHT metatypes.

A LIGHT component is an independently instantiable and deployable unit of computation and information that encapsulates reusable blocks of logic. Domain-specific type definitions for components specify compatible implementation and interface types. Examples of domain-specific component types include "JavaBean," "web server," or, in the case of the LL, "MyxComponent." Components define interaction points in the form of LIGHT interfaces. The domain-specific type definition for an interface restricts the allowed types of data and control exchange between components. Examples of domain-specific interface types include "HTTP port," "publish-subscribe API," or, in the case of the LL, "MyxAsynchronousInterface."

The relationship between an interface and a component has a designated *direction* (recall Figure 2), which can be either `implements` or `invokes`. The direction affects the semantics of the interface in several ways. First, the flow of information in the two directions is opposite: an invoking component sends the interface's inputs and receives its outputs; an implementing component receives the inputs and sends the outputs. Second, the *mode of interaction*—either *method-based* or *message-based*—is derived from the directions of the interfaces that are connected via a given link: method-based interactions occur when an invoked interface (*e.g.*, an object reference) is linked to an implemented interface (*e.g.*, an object), while message-based interactions occur when two invoked interfaces are linked (*e.g.*, two send-message/receive-message interfaces).

Resources are provided by the computing environment and are used by application implementations to perform tasks.

Under Review

	Embedded Semantic Assumptions		Properties	
	Capabilities	Responsibilities	Name (Type)	Description
`Component`	• Manage and prioritize interactions between internal implementations and external links • Multiplex/demultiplex, filter, and monitor interactions • Delegate externally initiated interactions to component implementations • Transmit internally initiated interactions to external entities via established links	• Specify mappings from implemented interfaces to subcomponent interfaces or tasks • Ensure interaction takes place via interfaces and protect internal information and behavior from being manipulated directly	`port` (Interface association)	Specifies the types of interfaces through which the component's implementations may interact
			`functionality` (Implementation association)	Designates the types of implementations that may be used to realize the component's provided services
			`composite` (Boolean attribute)	Indicates that the component is a hierarchy of subcomponents
`Interface`	• Ensure type conformance (data adheres to input and output data type definitions) • Ensure mode conformance (participants in an interaction are uniformly method- or message-based) • Ensure control conformance (an execution thread is transferred iff both source and target interfaces expect it) • Block an execution thread to create synchrony	• Declare at least one operation	`signature` (Operation association)	Indicates the operation types that may be specified by the interface type
			`control` (Boolean attribute)	Enables exchange of control flow across the interface type
			`synchronous` (Boolean attribute)	Forces synchronization across the interface; the interface invoker(s) and invokee(s) experience interactions simultaneously
			`queuing` (Boolean attribute)	Allows queuing of interactions over the interface
`Resource`	• Accept, queue, execute, and return service requests • Manage the allocation of pooled or divisible resources to requests • Maximize request execution parallelism as allowed by the metatype property specification	• Specify the available quantity/capacity • Specify a scheduling discipline	`class` (Enumeration attribute)	Specifies whether the resource is a computation (processing) resource, a communication resource, or an information (data) resource
			`units` (Enumeration attribute)	Defines whether the units of the resource are continuous or discrete
			`protocol` (Enumeration attribute)	Allows the resource to exhibit holding behavior, in which the resource is held until released by the service requester
			`instantiable` (Boolean attribute)	Permits the resource to be created on demand

Figure 3. Metatype semantics for the `Component`, `Interface`, and `Resource` metatypes.

Resources require (simulated) time in order to fulfill task requests, and contention over resources results in the emergent behavior of applications. Each resource may optionally permit an arbitrary level of parallelism; in other words, a resource may be capable of servicing multiple requests simultaneously. Resources are required to specify the available quantity of the resource (the *capacity*), which may be a positive real number if the resource is continuous or a natural number if the resource is discrete. For example, bandwidth may be modeled as a continuous resource, while threads in a thread pool are an example of a discrete resource. Resources may be processing (e.g., CPUs or threads), communication (e.g., network interfaces), or data resources (e.g., files or buffers).

To illustrate how customized semantics are specified in a LIGHT metamodel, we return to the Myx metamodel introduced in Section II. Myx allows for multiple types of interfaces with different semantics and constraints. The following semantic definitions are taken verbatim from the Myx specification (emphasis added):

1. The default form for an *interface* in Myx is a set of one or more *methods* that can be called.
2. Every interface is designated as either a *provided* or a *required* interface...
3. *Links* have exactly two endpoints. Each link connects exactly one required interface to one provided interface.
4. All bricks have two "domains" called *top* and *bottom*. All interfaces on a brick must be assigned to one of these domains.
5. In a *synchronous* invocation ... the calling component passes its thread of control to the called component, [which] completes its invocation and returns control to the calling component.
6. In an *asynchronous* invocation, the invoking brick continues processing after initiating the invocation, which proceeds concurrently in a separate thread of control.

The excerpt of the Myx metamodel shown in Figure 4 illustrates how the above semantics are captured in LIGHT. `MyxInterface`, an abstract base type for other interfaces, declares one or more `JavaMethodDecl` operations (capturing #1 from the table above). The `direction` of the `port` properties of `MyxBrick` is `implements` for `MyxProvidedInterface` and `invokes` for `MyxRequiredInterface` (#2). `MyxSyncLink` and `MyxAsynchLink` connect exactly one `MyxProvidedInterface` to exactly one `MyxRequiredInterface` (#3). `MyxBrick` has `topDomain` and `bottomDomain` properties, which are sets of `MyxInterfaces` (#4). A `MyxSynchronousInterface` passes a thread of control, mandates that the invoker and invokee experience interactions simultaneously, and disallows queueing of interactions (#5). A `MyxAsynchronousInterface`, on the other hand, does not pass a thread of control, allows the invoker to initiate an interaction that is experienced by the invokee at a later time, and allows interaction queueing (#6).

The metamodel for a different architecture DSL would have different types and semantics defined. For example, the metamodel of AADL [14], which is targeted at embedded real-time systems, includes execution platform components, such as processors and memory. AADL Interfaces are called *ports*. One type of port, the *event port*, causes an immediate transfer of control, but also can be queued at the recipient, resulting in different semantics than the interfaces found in Myx.

Using LIGHT, if engineers wish to change the semantics of their DSL, or add new DSL types, they can simply update the metamodel and regenerate the supporting tools. For example, if the LL development team decided that they wanted to alter the standard Myx semantics in some way (e.g., by making `MyxLinks` delay-free), they could do so by changing a single property, as opposed to dealing with low-level interpreter code.

Under Review

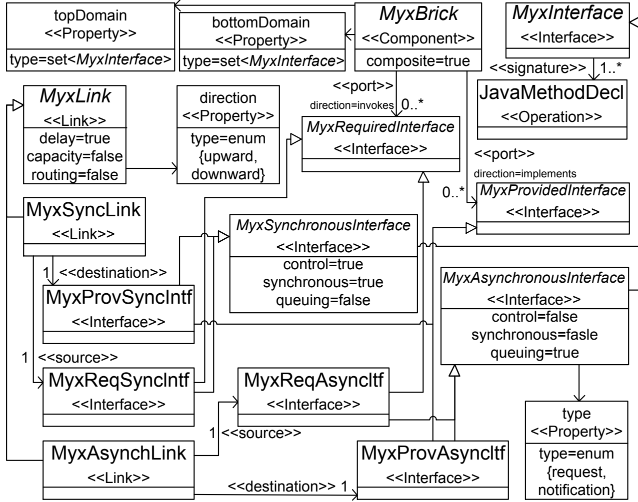

Figure 4. An excerpt of the Myx metamodel that defines `Interface` types.

Finally, since the semantic variability permitted by the LIGHT metalanguage is restricted by the set of defined semantic assumptions and properties for each metatype—meaning that some DSL semantics cannot be captured in a LIGHT metamodel—LIGHT provides a "backdoor" mechanism for defining semantics: engineers are able to insert custom behavioral definitions directly at the modeling level. LIGHT allows any default behavior to be overridden with a custom definition by providing a pointer to an external semantic specification file. For example, LIGHT does not provide a means for links to *correlate* interaction events, that is, to hold one event indefinitely until another event satisfying a condition occurs. However, such semantics can be added by specifying event correlating behavior directly in the native language of the target runtime platform. Custom semantics must be directly executable because LIGHT cannot translate this logic and treats custom behavior definitions as a black box.

B. Interpretation

Recall from Section III that domain-specific interpreters are composed of an MIF and autogenerated code for the MIF (an MIF extension). This section focuses on the implementation of MIFs and MIF extensions in LIGHT, with a particular emphasis on the reusable aspects of their design.

The architecture of LIGHT interpreters uses the *visitor-traverser* style [20]. In accordance with this style, each model interpreter is structured as two distinct modules: one to traverse models and another to generate code (see Figure 5). The model traversal module contains the logic for navigating through the model and invoking the code generator on each model object. The code generator module contains the logic for outputting code for the target analysis, simulation, or execution platform.

The use of the visitor-traverser style provides a basis for increasing the flexibility of MIFs by decoupling domain-independent interpretation logic (i.e., logic that does not vary from one metamodel to another) from domain-specific logic

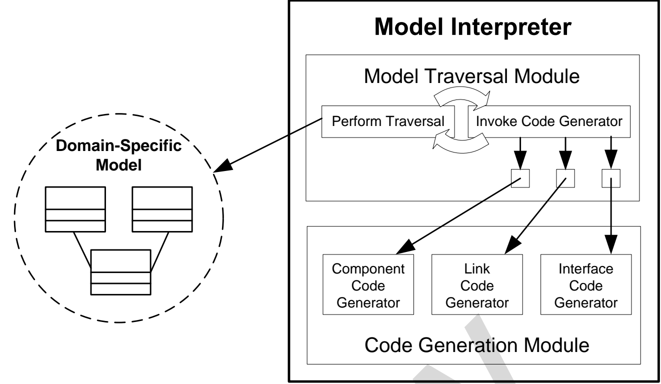

Figure 5. LIGHT model interpreter functions.

(i.e., logic that depends on the types and properties defined in a metamodel). This decoupling is necessary because classes implementing the domain-specific logic must be individually generated and customized for each metamodel.

The code generation module of an LIGHT model interpreter is a combination of classes built into the MIF and autogenerated classes contained in an MIF extension. Returning to Figure 1, the built-in, domain-independent classes are represented by the *Transformation Logic* within the MIF; these classes implement code generation operations. The autogenerated, domain-specific classes are represented by the *Transformation Rules*; these classes contain instructions for applying code generation operations.

At the implementation level, the domain-independent logic is contained within MIF classes corresponding to each metatype, called *metatype classes*, while the domain-specific logic is contained within MIF extension classes corresponding to each metatype instance (i.e., each domain-specific type defined in a metamodel), called *type classes*. The metatype classes define *template methods* that implement specific interpretation tasks, and these methods are selectively invoked and parameterized by the type classes. In this way, the MIF implements the high-level structure of interpretation algorithms but allows customization of the algorithms by MIF extensions.

To illustrate how semantics specified in a LIGHT metamodel are encoded as transformation rules within an MIF extension, we return to the Myx version of the LL model. Recall that, in the Myx architectural framework, component interfaces may be synchronous or asynchronous. Synchronous interfaces have the same semantics as Java method calls, and accordingly the following properties are specified in the Myx metamodel for the `MyxSynchronousInterface` type: `control=true`, `synchronous=true`, `queuing=false`. Consequently, the Myx MIF extension for LIGHT's simulation generator MIF contains a `MyxSynchronousInterface` type class, which invokes the `generateThreadPasser()` method of the `Interface` metatype class, but does not invoke its `generateEventQueue()` method.

The described design results in several important benefits:

Under Review

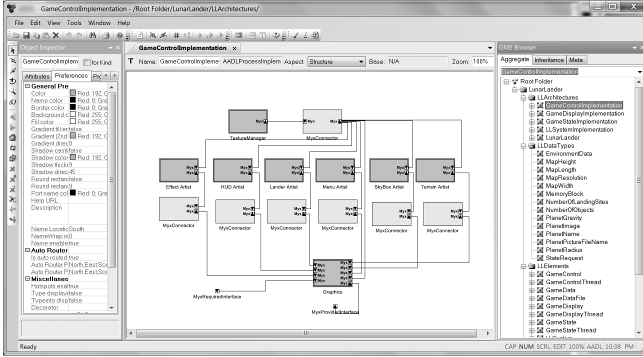

Figure 6. A screenshot of an LL model in the automatically generated Myx model editor.

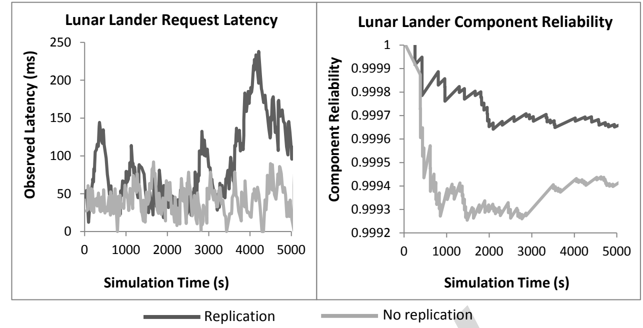

Figure 7. Latency and error rates computed in a Lunar Lander simulation.

- The design factors out the common structure of model interpretation algorithms to simplify code reuse, while still providing a straightforward way for domain-specific logic to customize the interpretation process.
- The design allows the MIF to limit and control the ways in which interpretation may be customized, ensuring that a MIF extension does not "break" the interpretation process or violate the constraints of the target runtime platform.
- The design ensures consistency of metatype properties and avoids duplication of common properties by allowing metatype properties and domain-specific properties to be separated within distinct classes.

C. Simulation

LIGHT automatically generates fully configured, domain-specific interpreters that today have to be programmed manually. One type of interpreter generated by LIGHT is simulation generators for XDEVS, a stand-alone simulator for analyzing the dynamic behavior of complex systems [11]. XDEVS is a component-based variant of the widely used Adevs [27] event-based simulation platform. LIGHT's simulation generators consist of the XDEVS MIF (built into LIGHT) and domain-specific XDEVS MIF extension code (autogenerated by LIGHT). We have used the LIGHT-generated XDEVS simulators to perform latency [39], memory usage, energy consumption [34], and reliability [30] analyses. Here, we briefly describe how simulation can be used to analyze the Myx LL.

Suppose an engineer wishes to evaluate how replication of a critical LL component might affect performance and reliability. While running an additional replica will likely require more resources and decrease performance, the system reliability will improve. Comparing the two possible designs is difficult because performance and reliability depend on the complex interactions of numerous components, the execution environment, and other factors.

Figure 6 shows a screenshot of the automatically generated Myx model editor with one of the two possible LL designs. We used LIGHT to generate XDEVS simulation code from both LL designs, modeled in Myx. We used XDEVS's native facilities to instrument the simulation with probes and quantify the two properties of interest. Figure 7 summarizes the end-to-end latency and failure rate data from the simulations. The two graphs quantify the intuition stated above. The graphs allow an engineer to make an informed architectural decision and provide quantitative rationale regarding whether the performance cost of replication is acceptable, or whether the system reliability requirements can be met without replication.

V. Evaluation

To evaluate the utility of LIGHT, we created metamodels for nine software architecture DSLs. We simulated domain-specific system models conforming to these metamodels using LIGHT's XDEVS simulation generator components and the XDEVS discrete event simulation engine.[1] Note that the ability to generate such simulations from architectural models is not novel; rather, the novelty comes from the ability to do so without writing any custom code by hand. Therefore, we evaluated LIGHT according to programming effort it saved. In this section, we define our evaluation metrics and present the results of applying LIGHT to the nine DSLs.

A. Metrics

We used two sets of evaluation metrics: *Implementation effort* metrics measure the effort saved through code generation and reuse. *Maintenance effort* metrics measure the relative ease of performing DSL modifications in a metamodel, as afforded by LIGHT, rather than in interpreter source code, as is necessary in other MDE platforms.

The implementation effort metrics are:

- *Number of domain-specific types* in the metamodel, which is a measure of metamodel size and complexity.
- *Lines of generated interpreter code* by the metainterpreter in the form of MIF extensions.
- *Total lines of reused interpreter code*, which is the sum of (1) the 16,243 SLOC in the XDEVS MIF (built into LIGHT for reuse) and (2) the generated SLOC in the MIF extension for each metamodel.
- *Lines of generated code per domain-specific type.*
- *Lines of reused code per domain-specific type.*

Under Review

Metamodel	Description
AADL	Modeling language targeted for specification and analysis of real-time embedded systems.
xADL 2.0 Core	A common set of fundamental modeling abstractions for software architectures.
Ecore	Eclipse metamodel for object-oriented systems.
C2	Component- and message-based architectural style for flexible, extensible software systems.
Client/server	Ubiquitous request/response architectural style.
Pipe-and-filter	Concurrent data-stream architectural style.
Publish-subscribe	Asynchronous and anonymous message distribution architectural style.
Myx	Layered architectural style for flexible construction of distributed systems.
Prism	Event-based style for embedded, mobile applications.

Figure 8. Nine metamodels used in the evaluation.

metamodel	domain-specific types	sloc generated	total sloc reuse	generated sloc per type	total reuse per type
AADL	86	2,306	18,549	27	216
Ecore	76	1,277	17,520	17	231
xADL Core	53	1,195	17,438	23	329
C2	42	1,268	17,511	30	417
Client/server	34	1,091	17,334	32	510
Pipe-and-filter	29	957	17,200	33	593
Pub-sub	50	1,779	18,022	36	360
Myx	62	1,493	17,736	24	286
Prism	54	1,678	17,921	31	332

(a)

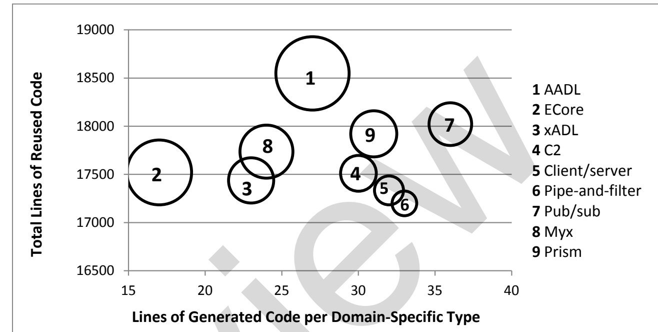

(b)

Figure 9. The collected implementation effort metrics (a) are summarized based on the metamodel size, total reused code, and generated code per type (b).

The maintenance effort metrics were computed by (1) performing a set of modifications to a DSL, (2) regenerating the XDEVS MIF extension for the DSL, and (3) diffing the previous and new MIF extensions to determine the impact of the DSL changes. The maintenance effort metrics are:

- *The number of metamodel objects altered.*
- *The number of interpreter classes altered.*
- *The number of interpreter methods altered.*
- *Lines of code altered per domain-specific type.*
- *Total lines of code altered.*

The first metric provides an indication of the level of effort needed to achieve a modification using the LIGHT approach; the remaining metrics indicate the level of effort that would be required using a conventional approach.

B. Results

Figure 8 summarizes our nine subject metamodels. For each metamodel, Figure 9a displays the implementation effort metrics for the XDEVS simulation generator components. (The model editor components and code generated for a specific application model, such as LL, would also be automatically generated using previous approaches, so these metrics are not included.) Figure 9a shows that the DSLs range in the number of domain-specific types defined (a rough indicator of the metamodeling effort required). This is to be expected as the nine DSLs vary widely in their scope and features. Using a conventional approach, the number of domain-specific types would be the same, but each domain-specific type would be somewhat easier to define because there are fewer metatype property values to set. This is the single added cost of using our approach as compared to existing techniques; however, this cost is eclipsed by the resulting savings, as discussed below.

Figure 9b illustrates the relationship between metamodel size and two metrics — generated code per type and total code reuse. The diameter of each circle represents the size of the corresponding metamodel. The larger metamodels generally result in greater overall reuse, but the reuse benefits are amortized over a larger metamodeling effort, resulting in a smaller benefit per domain-specific type. This reinforces the intuition that, for very large metamodels, more implementation effort is avoided, but metamodeling consumes a larger share of the overall effort.

To estimate the savings incurred by using LIGHT, we applied the widely-used COCOMO II model [1]. As indicated in Figure 9a, the average amount of code reused in the implementation of domain-specific simulation generators for the nine DSLs is $\sim$17,500 SLOC. Let us assume that this amount of code had to be written manually.[2] COCOMO II's estimates of the required effort for a project of this size range between 4.2 and 23.4 person-months, depending on the project parameters.

Figure 10 shows the maintenance effort metrics for seven DSL modifications. To obtain a broad sample, we modified different metamodels in different ways, including introducing new inheritance and containment relationships, changing metamodel properties, and adding or removing types. The resulting changes in the generated code depended on the specific modifications. For example, inheritance modifications tended to affect a large number of classes but may not have changed any methods within those classes (as in the case of client/server). In contrast, containment modifications affected a small number of classes but resulted in more new code (as in the case of AADL).

While the number of lines of code altered in response to the changes was, in some cases, quite small, these changes tended to be spread over a relatively large number of classes

[1] We also generated executable Prism-MW code. The results were comparable to those we show for simulation and we omit them for space.

[2] While we do not have access to handwritten versions of these simulation generators, we expect that their size would be proportional to the size of the automatically generated versions.

Under Review

metamodel	types affected	classes altered	methods altered	sloc altered per type	sloc altered
AADL	2	3	10	64.5	129
Client/server	4	10	1	5	20
C2	2	4	1	7	14
Ecore	1	4	4	19	19
Pub-sub	1	3	4	22	22
Myx	5	8	12	20.4	102
Prism	4	7	14	21	84

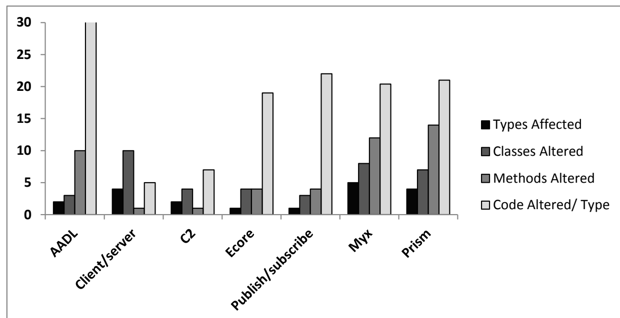

Figure 10. The collected maintenance effort metrics.

and methods in comparison to the relatively small number of metamodel objects that were altered. This implies that modifications required for DSL evolution are more widely scattered, and therefore more time-consuming to implement, in interpreter source code than in a metamodel. Moreover, a metamodel may be more understandable and easier for a new engineer to modify, which is important when the DSL must be maintained over a period of time.

VI. Related Work

Our work builds upon prior research in architecture modeling [6], [25], [35], product line architectures [3], [17], [40], model-driven engineering [21], [33], [38], model-driven architecture [28], [29], [36], system analysis and simulation [2], [4], [37], and application frameworks [13], [32]. Here, we specifically highlight three notable DSL-based approaches that attempt to automate generation of model interpreters.

Cadena [18] is a model-driven toolchain for component-based systems. Bogor is an extensible model checking framework that can be applied to Cadena models. Similarly to LIGHT, Bogor allows engineers to reuse the analysis infrastructure while customizing selected constructs and behaviors. However, some of the mechanisms that enable extension and reuse in Bogor are specific to model checking, making them difficult to generalize and apply to other types of analysis or code generation. In this paper, we have attempted to define a common, reusable architecture that can be applied in a variety of contexts.

The DUALLY framework [23] supports interoperability among architecture DSLs. DUALLY leverages engineer-defined mappings between DSL metamodels and a core set of architectural concepts codified in a metamodel, called A_0. Thus, the analysis and code generation tools available for any DSL with a defined mapping can be applied to architectural models specified in other DSLs. Although DUALLY eliminates the need to manually program model transformations, it still requires defining the mappings between languages. DUALLY also exacerbates the traceability problem: relating analysis results or bugs in generated code to the source architectural model. Introducing additional model transformations makes traceability more challenging.

The ALFAMA workbench [31] automates the construction of DSLs for application frameworks. ALFAMA leverages an aspect-oriented domain-specific modeling (DSM) layer that defines *specialization aspects* that modularize framework hot-spots (extension points) and associates those aspects with DSL concepts. The DSM layer allows engineers to bypass metamodel development, resulting in an approach that is, in some ways, the reverse of ours: rather than enhancing metamodels to eliminate interpreter development, ALFAMA enhances interpreters to eliminate metamodel development. This has many merits but also some drawbacks. First, high-level metamodels are more maintainable than low-level source code. Second, inferring DSLs from framework hot-spots tightly couples the DSL to a particular framework. As with LIGHT, the automation gains of the ALFAMA approach reduce DSL flexibility in some circumstances.

VII. Contributions

The attraction of being able to automatically synthesize model editing, analysis, and code generation tools is that it greatly reduces the cost of using DSLs. In turn, this makes the benefits of domain-specific modeling, traditionally enjoyed by large-scale defense and aerospace programs, possible for small- and medium-sized software projects. Our approach opens this possibility: we exploit an isomorphism between model editors and other model interpreters to automatically synthesize interpreters from DSL metamodel definitions.

We implemented our approach in the LIGHT platform, targeted at architecture-based DSLs. We evaluated LIGHT by constructing three different model interpreters across nine DSLs to demonstrate that it significantly reduces the required programming effort and system maintenance, and aids design exploration.

LIGHT does not come without limitations. One seeming limitation is a slightly increased metamodeling burden over existing MDE approaches. However, this is more than compensated for by the implementation savings. Another, real limitation stems from the trade-off between the expressiveness and flexibility of the supported DSLs on the one hand, and the ability to automatically synthesize tools on the other. Precisely qualifying, and possibly quantifying, this trade-off is part of our ongoing work.

Under Review

References

[1] B. Boehm *et al.*, *Software Cost Estimation with Cocomo II*. Prentice Hall, 2000.

[2] E. M. Clarke and B.-H. Schlingloff, "Model Checking," in *Proceedings of the 17th Conference on Foundations of Software Technology and Theoretical Computer Science*, Kharagpur, India, December 1997, pp. 54–56.

[3] P. Clements and L. Northrop, *Software Product Lines*. Addison-Wesley Professional, 2001.

[4] J. B. Dabney and T. L. Harman, *Mastering Simulink*. Prentice Hall, 2003.

[5] E. Dashofy, "Supporting Stakeholder-driven, Multi-view Software Architecture Modeling," Ph.D. dissertation, UCI, 2007.

[6] E. M. Dashofy, A. van der Hoek, and R. N. Taylor, "An Infrastructure for the Rapid Development of XML-based Architecture Description Languages," in *Proceedings of the 24th International Conference on Software Engineering (ICSE)*, Orlando, Florida, May 2002, pp. 266–276.

[7] G. Edwards *et al.*, "Construction of Analytic Frameworks for Component-Based Architectures," in *SBCARS*, 2007.

[8] G. Edwards, S. Malek *et al.*, "Scenario-Driven Dynamic Analysis of Distributed Architectures," in *FASE*, 2007.

[9] G. Edwards and N. Medvidovic, "A Methodology and Framework for Creating Domain-Specific Development Infrastructures," in *ASE*, 2008.

[10] G. Edwards, Y. Brun, and N. Medvidovic, "Isomorphism in model tools and editors," in *ASE*, November 2011.

[11] G. Edwards *et al.*, "A Highly Extensible Simulation Framework for Domain-Specific Architectures," University of Southern California, Center for Software Engineering, Tech. Rep. USC-CSSE-2009-511, 2009.

[12] G. Edwards, "Automated Synthesis of Domain-Specific Model Interpreters," Ph.D. dissertation, USC, 2010.

[13] M. Fayad and D. C. Schmidt, "Object-Oriented Application Frameworks," *Communications of the ACM*, vol. 40, no. 10, pp. 32–38, October 1997.

[14] P. H. Feiler, D. P. Gluch, and J. J. Hudak, "The Architecture Analysis & Design Language: An Introduction," Software Engineering Institute, Carnegie Mellon University, Tech. Rep. CMU/SEI-2006-TN-011, February 2006.

[15] "The Generic Modeling Environment," http://www.isis.vanderbilt.edu/Projects/gme/.

[16] "The Eclipse Graphical Modeling Framework," http://www.eclipse.org/modeling/gmf/.

[17] S. A. Hendrickson and A. van der Hoek, "Modeling Product Line Architectures through Change Sets and Relationships," in *Proceedings of the 29th International Conference on Software Engineering (ICSE)*, Minneapolis, Minnesota, May 2007, pp. 189–198.

[18] G. Jung *et al.*, "A Type-centric Framework for Specifying Heterogeneous, Large-scale, Component-oriented, Architectures," in *GPCE*. ACM, 2007, p. 42.

[19] G. Karsai *et al.*, "On the Use of Graph Transformation in the Formal Specification of Model Interpreters," *J. Universal Computer Science*, vol. 9, no. 11, pp. 1296–1321, 2003.

[20] G. Karsai, "Structured Specification of Model Interpreters," in *ECBS*, Nashville, Tennessee, March 1999, pp. 84–90.

[21] S. Kelly, "Towards a Comprehensive MetaCASE and CAME Environment: Conceptual, Architectural, Functional and Usability Advances in MetaEdit+," Ph.D. dissertation, University of Jyvaskyla, Jyvaskyla, Finland, 1997.

[22] "LIGHT Metatype Quick Reference Guide," http://softarch.usc.edu/~gedwards/xteam2/MetatypeTable.pdf.

[23] I. Malavolta *et al.*, "Providing Architectural Languages and Tools Interoperability through Model Transformation Technologies," *IEEE Transactions on Software Engineering*, 2009.

[24] S. Malek *et al.*, "A Style-Aware Architectural Middleware for Resource-Constrained, Distributed Systems," *IEEE TSE*, vol. 31, no. 3, pp. 256–272, March 2005.

[25] N. Medvidovic and R. N. Taylor, "A Classification and Comparison Framework for Software Architecture Description Languages," *IEEE Transactions on Software Engineering*, vol. 26, no. 1, pp. 70–93, January 2000.

[26] P. Mohagheghi and V. Dehlen, "Where Is the Proof? - A Review of Experiences from Applying MDE in Industry," in *Proceedings of the 4th European conference on Model Driven Architecture Foundations and Applications (ECMDA-FA)*, Berlin, Germany, June 2008, pp. 432–443.

[27] J. J. Nutaro, *Building Software for Simulation: Theory and Algorithms, with Applications in C++*. Wiley Publishing, 2010.

[28] *MDA Guide, V1.0.1*, Object Management Group, June 2003.

[29] *Meta Object Facility (MOF) Core Specification, V2.0*, Object Management Group, January 2006.

[30] R. Roshandel, S. Banerjee, L. Cheung, N. Medvidovic, and L. Golubchik, "Estimating Software Component Reliability by Leveraging Architectural Models," in *Proceedings of the 28th International Conference on Software Engineering (ICSE)*, Shanghai, China, May 2006, pp. 853–856.

[31] A. L. Santos *et al.*, "Automating the Construction of Domain-Specific Modeling Languages for Object-Oriented Frameworks," *Journal of Software and Systems*, 2010.

[32] A. L. Santos, A. Lopes, and K. Koskimies, "Framework Specialization Aspects," in *Proceedings of the 6th International Conference on Aspect-Oriented Software Development (AOSD)*, Vancouver, Canada, March 2007, pp. 14–24.

[33] D. C. Schmidt, "Guest Editor's Introduction: Model-Driven Engineering," *IEEE Computer*, vol. 39, no. 2, pp. 25–31, February 2006.

[34] C. Seo, "Prediction of Energy Consumption Behavior in Component-Based Distributed Systems," Ph.D. dissertation, University of Southern California, Los Angeles, California, 2008.

[35] R. N. Taylor *et al.*, *Software Architecture: Foundations, Theory and Practice*. Wiley Publishing, 2009.

[36] "Unified Modeling Language," http://www.uml.org.

[37] A. Varga, "The OMNeT++ Discrete Event Simulation System," in *Proceedings of the European Simulation Multiconference (ESM)*, Prague, Czech Republic, June 2001, pp. 319–324.

[38] K. C. Wallnau, "Volume III: A Technology for Predictable Assembly from Certifiable Components," Software Engineering Institute, Carnegie Mellon University, Tech. Rep. CMU/SEI-2003-TR-009, April 2003.

[39] M. Woodside, "Tutorial Introduction to Layered Modeling of Software Performance," Department of Systems and Computer Engineering, Carleton University, Tech. Rep., May 2002.

[40] H. Ye and H. Liu, "Approach to Modelling Feature Variability and Dependencies in Software Product Lines," *IEE Proceedings - Software*, vol. 152, no. 3, pp. 101–109, June 2005.

Under Review