

Component-based System Integration via (Meta)Model Composition*

Krishnakumar Balasubramanian, Douglas C. Schmidt, Zoltán Molnár, Ákos Lédeczi
Institute for Software Integrated Systems
Vanderbilt University, Nashville

{kitty,schmidt,zolmol,akos}@isis.vanderbilt.edu

Abstract

This paper provides three contributions to the study of functional integration of distributed enterprise systems. First, we describe the challenges associated with functionally integrating the software of these systems. Second, we describe how the composition of domain-specific modeling languages (DSMLs) can simplify the functional integration of enterprise distributed systems by enabling the combination of diverse middleware technologies. Third, we demonstrate how composing DSMLs can solve functional integration problems by reverse engineering an existing CORBA Component Model (CCM) system and exposing it as Web Service(s) to web clients who use these services. This paper shows that functional integration done using (meta)model composition provides significant benefits with respect to automation, reusability, and scalability compared to conventional integration processes and methods.

1. Introduction

1.1. Challenges of Functional Integration

With the emergence of commercial-off-the-shelf (COTS) component middleware technologies, such as Enterprise Java Beans (EJB), CORBA Component Model (CCM), and Microsoft .NET Framework, software developers are increasingly faced with the task of integrating heterogeneous enterprise distributed systems built using different COTS technologies, rather than just integrating proprietary software developed in-house. Although there are well-documented patterns [8] and techniques [4] for system integration using various component middleware technologies, system integration is still largely a tedious and error-prone manual process. To improve this process, component developers and system integrators must therefore understand key properties of the

systems (“system” refers to an enterprise distributed system built using component middleware like EJB, Microsoft .NET, or CCM in the remainder of this paper) they are integrating, as well as the integration technologies they are applying.

This paper describes technologies that help simplify the *functional integration* of systems built using component middleware. This type of integration operates at the logical business layer, typically using distributed objects/components, exposing existing functionality as services, or using messaging middleware. Functional integration of systems is hard due to the variety of available component middleware technologies. These technologies differ in many ways, including the protocol level, the data format level, the implementation language level, and/or the deployment environment level. In general, however, component middleware technologies are a more effective technology base than the brittle proprietary infrastructure used in legacy systems, which have historically been built in a vertical, stove-piped fashion.

Despite the benefits of component middleware, key challenges in functional integration remain unresolved when integrating large-scale systems developed using heterogeneous COTS middleware. These challenges include (1) *integration design*, which involves choosing the right abstraction for integration, (2) *interface mapping*, which reconciles different datatypes, (3) *technology mapping*, which reconciles various low-level issues, (4) *deployment mapping*, which involves planning the deployment of heterogeneous COTS middleware, and (5) *portability incompatibilities* between different implementations of the same middleware technology. The lack of simplification and automation in resolving these challenges significantly hinders effective system integration.

1.2. Solution Approach→Functional Integration using (Meta)Model Composition

A promising approach to address the integration challenges outlined above is *Model-Driven Engineering* (MDE), which involves the systematic use of models as essential artifacts throughout the software lifecycle [20]. At

* This work was sponsored in part by grants from Raytheon and Lockheed Martin Advanced Technology Laboratories

the core of MDE is the concept of *domain-specific modeling languages* (DSMLs) [12], whose type systems formalize the application structure, behavior, and requirements within particular domains, such as software defined radios, avionics mission computing, online financial services, warehouse management, or even the domain of middle-ware platforms.

While DSMLs have been used to help software developers create homogeneous systems [11, 21], enterprise distributed systems are rarely homogeneous. A single DSML developed for a particular component middleware technology, such as EJB or CCM, may therefore not be applicable to model, analyze, and synthesize key concepts of Web Services. To integrate heterogeneous systems successfully, therefore, system integrators need tools that can provide them with a unified view of the entire enterprise system, while also allowing them fine-grained control over specific subsystems and components.

Our approach to integrating heterogeneous systems is (*meta*)*model composition*¹ [1], which (1) creates a new DSML from multiple existing DSMLs by adding new elements or extending elements of existing DSMLs, (2) specifies new relationships between existing elements, and (3) defines relationships between new and existing elements. This paper describes *System Integration Modeling Language* (SIML), which is our open-source DSML that enables functional integration of component-based systems via (*meta*)*model composition*. We developed SIML using the Generic Modeling Environment (GME) [13], which is an open-source meta-programmable modeling environment.

2. Functional Integration of Component-based Systems

To motivate the need for model-driven functional integration capabilities, this section describes the challenges in functional integration of component-based systems.

2.1. Functional Integration Challenges

Functional integration of systems is hard and involves activities that map between various levels of abstraction in the integration lifecycle, including design, implementation, and use of tools. We now describe key challenges associated with integrating older component middleware technologies, such as CCM and EJB, with newer technologies, such as Web Services.

Challenge 1. Choosing an appropriate level of integration. As shown in Step 1 of Figure 1, a key activity is to

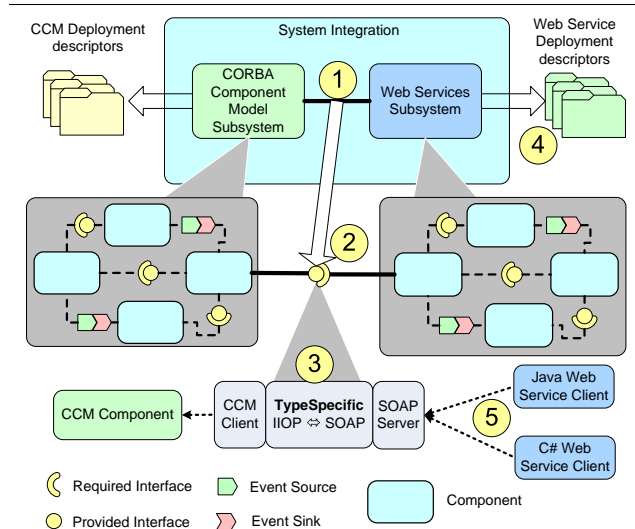


Figure 1. Functional Integration Challenges

identify the right level of abstraction at which functional integration should occur, which involves selecting elements from different technologies being integrated that can serve as conduits for exchanging information. Among the different possible levels at which integration can be performed, criteria for determining the appropriate level of integration include: (1) the number of normalizations, *i.e.*, conversion to/from the native types, required to ensure communication between the peer entities being integrated, (2) the number (and hence the overhead) of as well as the flexibility of deployment, *i.e.*, in-process/out-of-process etc. of runtime entities required to support the functional integration, (3) the number of changes to the integration architecture required corresponding to changes to the peers that are being integrated, (4) available choices of platform-specific infrastructure, *i.e.*, operating systems, programming languages *et al.*, associated with performing integration at a particular level. Attempting integration at the wrong level of abstraction can yield brittle integration architectures that break when changes occur to either the source or target system being integrated.

Challenge 2. Reconciling differences in interface specifications. After the level of abstraction to perform functional integration is determined, it is necessary to map the interfaces exposed by elements of the different technologies as shown in Step 2 of Figure 1. Common COTS middleware technologies usually have an interface definition mechanism that is separate from the implementation details, *e.g.* CCM uses the OMG Interface Definition Language (IDL), whereas Web Services use W3C Web Services Definition Language (WSDL). Irrespective of the mechanism used to define interfaces, mapping of interfaces between any two

¹ The term “(meta)model” conveys the fact that this composition technique can be applied to both metamodels and models.

technologies involves at least three tasks: (1) *datatype mapping*, which involves mapping a datatype (both pre-defined and complex types) from source to target technology, (2) *exception mapping*, which involves mapping exceptions from source to target technology; exceptions are not clubbed together with datatypes since the source or target technologies might not have a notion of exceptions (e.g. Microsoft's COM uses a *HRESULT* to convey errors instead of using exceptions), and (3) *language mapping*, which involves mapping datatypes between two technologies while accounting for differences in languages at the same time.²

Performing these mappings is non-trivial, requires expertise in both the source and target technologies, and exposes severe scalability problems due to their tedium and error-proneness if they are not automated.

Challenge 3. Managing differences in implementation technologies. The interface mapping described above addresses the high-level details of how information is exchanged between different technologies being integrated. As shown in Step 3 of Figure 1, however, low-level technology details such as networking, authentication and authorization *et al.* are responsible to actually delivering such integration. This involves a technology mapping and includes the following activities: (1) *protocol mapping*, which reconciles the differences between the protocols used for communication between the two technologies, (2) *discovery mapping*, which allows bootstrapping and discovery of components/services between source and target technologies, and (3) *Quality of Service (QoS) mapping*, which maps QoS mechanisms between source and target technologies to ensure that service-level agreements (SLAs) are maintained.

Mapping of protocol, discovery, and QoS technology details requires not only expertise in the source/target technologies, but also intimate knowledge of the implementation details of these technologies.

Challenge 4. Managing deployment of subsystems. Component middleware technologies use declarative notations (such as XML descriptors, source-code attributes, and annotations) to capture various configuration options. Example metadata include EJB deployment descriptors, .NET assembly manifests, and CCM deployment descriptors. As shown in Step 4 of Figure 1, system integrators must track and configure metadata correctly during integration and deployment. In many cases, the correct functionality of the integrated system depends on correct configuration of the metadata.

Challenge 5. Dealing with interoperability issues. Unless a middleware technology has only one version implemented by one provider, there may be multiple implemen-

tations from different providers. As shown in Step 5 of Figure 1, differences between these implementations will likely arise due to non-conformant extension to standards, different interpretations of the same (often vague) specification, or implementation bugs. Regardless of the reasons for incompatibility, however, problems arise that often manifest themselves only during system integration. Examples of such differences are highlighted by the presence of efforts like the Web Services-Interoperability Basic Profile (WS-I) [3], which is a standard aimed at ensuring compatibility between the Web Services implementations from different vendors.

Due to the challenges described above, significant integration effort is spent on configuration activities, such as modifying deployment descriptors, and interoperability activities, such as handcrafting protocol adapters to link different systems together, which does not scale up as the number of components in the system increases or the number of adaptations required increases. Problems discovered at integration stage often require changes to the implementation, and thus necessitate interactions between developers and integrators. These interactions are often inconvenient, and even infeasible (especially when using COTS products), and can significantly complicate integration efforts. The remainder of this paper shows how our GME-based (meta)model composition framework and associated tools help address these challenges.

3. DSML Composition using GME

This section describes the (meta)model composition framework in the Generic Modeling Environment (GME) [13]. GME is a meta-programmable modeling environment with a general-purpose editing engine, separate view-controller GUI, and a configurable persistence engine. Since GME is *meta-programmable*, it can be used to design DSMLs, as well as build models that conform to a DSML.

DSMLs are defined by metamodels, hence, DSML composition is defined by (meta)model composition. The specification of how metamodels should be composed, *i.e.*, what concepts in the metamodels that are composed relate to each other and how, can be specified via normal association relationships and additional composition operators, as described in GME [1].

A key property of a composite DSML is that it supports the *open-closed* principle [15], which states that a class should be open for extension but closed with respect to its public interface. In GME, elements of the sub-DSMLs are *closed, i.e.*, their semantics cannot be altered in the composite DSML. The composite DSML itself, however, is *open, i.e.*, it allows the definition of new interactions and the creation of new derived elements. All tools that are

² Functional integration is very limited when attempting the latter mapping, which is often done via inter-process communication.

built for each sub-DSML work without any modifications in the composite DSML and all the models built in the sub-DSMLs are also usable in the composite DSML.

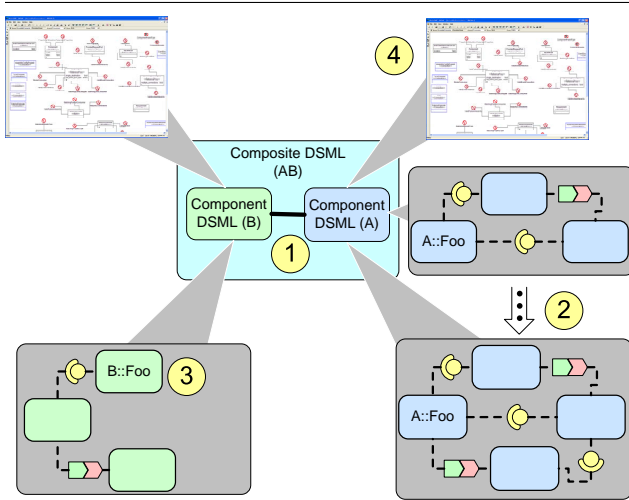


Figure 2. Domain-Specific Modeling Language Composition in GME

We use the following GME (meta)model composition features to support the SIML-based integration of systems built using different middleware technologies, as described in Section 4:

- **Representation of independent concepts.** To enable complete reuse of models and tools of the sub-DSMLs, the composition must be done in such a way that all concepts defined in the sub-DSMLs are preserved. As shown in Step 1 of Figure 2, no elements from either sub-DSMLs should be merged together in the composite DSML. GME's composition operators [1] can be used to create new elements in the composite DSML, but the sub-DSMLs as a whole must remain untouched. As a consequence, any model in a sub-DSML can be imported into the composite language, and vice versa. All models in the composite language that are using concepts from the sub-DSMLs can thus be imported back into the sub-DSML. Existing tools for sub-DSMLs can be reused as well in the composite environment. This technique of composing DSMLs is referred to as *metamodel interfacing* [6] since we create new elements and relationships that provide the interface between the sub-DSMLs.

- **Supporting (meta)model evolution.** DSML composition enables reuse of previously defined (sub-)DSMLs. Just like code reuse in software development, (meta)model reuse can also benefit from the concept of libraries. If an existing (meta)model is simply copied into new composite (meta)models, any changes or upgrades to the original will

not propagate to the places where they are used. As shown in Step 2 of Figure 2, if the original (meta)model is imported as a library, GME provides seamless support to update it when new versions become available (libraries are supported in any DSML with GME, not just the metamodeling language) Libraries are read-only projects imported to a host project. Components in the host project can create references to and derivations of library components. The library import process creates a copy of the reused project, so subsequent modifications to the original project are not updated automatically. To update a library inside a host project, a user-initiated refresh operation is required. To achieve unambiguous synchronization, elements inside a project have unique ids, which facilitates correct restoration of all relationships that are established among host project components and the library elements.

- **Partitioning (meta)model namespaces.** When two or more (meta)models are composed, name clashes may occur. To alleviate this problem, (meta)model libraries (and hence the corresponding components DSMLs) can have their own namespaces specified by (meta)modelers, as shown in Step 3 of Figure 2. External software components, such as code generators or model analysis tools that were developed for the composite DSML, must use the fully qualified names. But tools that were developed for component DSMLs will still work because GME sets the context correctly before invoking such a component.

- **Handling constraints.** The syntactic definitions of a metamodel in GME can be augmented by static semantics specifications in the form of Object Constraint Language (OCL) constraint expressions. When metamodels are composed together, the predefined OCL expressions coming from a sub-DSML should not be altered. Therefore GME's Constraint Manager uses namespace specifications to avoid any possible ambiguities, and these expressions are evaluated by the Constraint Manager with the correct types and priorities as defined by the sub-DSML as shown in Step 4 of Figure 2. The composite DSML can also define new OCL expressions to specify the static semantics that augment the specifications originating in the metamodels of the sub-DSMLs.

4. Integrating Systems with SIML

This section describes how we created and applied the *System Integration Modeling Language* (SIML), which is our open-source composite DSML that simplifies functional integration of component-based systems built using heterogeneous middleware technologies.

4.1. The Design and Functionality of SIML

Applying GME’s (meta)model composition features to SIML. To support integration of systems built using different middleware technologies, SIML uses the GME (meta)model composition features described in Section 3. SIML is thus a composite DSML that allows integration of systems by composing multiple DSMLs, each representing a different middleware technology. Each sub-DSML is responsible for managing the metadata (creation, as well as generation) of the middleware technology it represents. The composite DSML defines the semantics of the integration, which might include reconciling differences between the diverse technologies, as well as representing characteristics of various implementations. System integrators therefore have a single environment that allows the creation and specification of elements in each sub-DSML, as well as interconnecting them as if they were elements of a single domain.

Applying SIML to compose CCM and Web Services. Our initial use of SIML was to help integrate CCM with Web Services in the context of the shipboard computing case study described in Section 2. The two sub-DSMLs we needed to integrate to support the new requirements described in Section 2 were:

- The **Platform-Independent Component Modeling Language (PICML)**, which enables developers of CCM-based systems to define application interfaces, QoS parameters, and system software building rules, as well as generate valid XML descriptor files that enable automated system deployment.
- The **Web Services Modeling Language (WSML)**, which enables development of Web Services, and supports key activities in Web Service development, such as creating a model of a Web Services from existing WSDL files, specifying details of a Web Service including defining new bindings, and auto-generating artifacts required for Web Service deployment.

Since SIML is a composite DSML, all valid elements and interactions from both PICML and WSML are valid in SIML. It is therefore possible to design both CCM components (and assemblies of components), as well as Web Services (and federations of Web Services) using SIML, just as if either PICML or WSML were used independently. The whole is greater than the sum of its parts, however, because SIML defines new interactions that allow connecting a CCM component (or assembly) with a Web Service and automates generation of necessary gateways, which are capabilities that exist in neither PICML nor WSML.

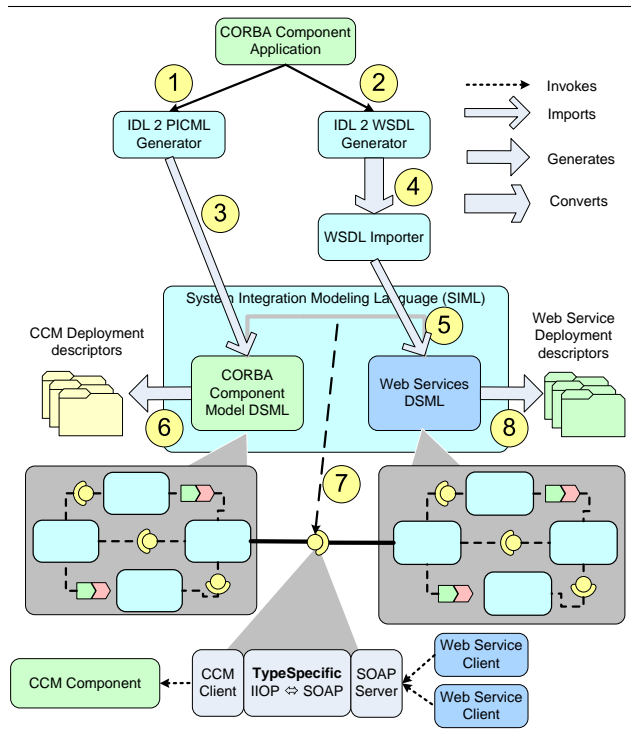


Figure 3. Generating a Web Service Gateway Using SIML

4.2. Resolving Functional Integration Challenges using SIML

We now show how we applied SIML to resolve the functional integration challenges discussed in Section 2.1. Although we focus on the initial version of SIML that supports integration of CCM and Web Services, its design is sufficiently general that it can be applied to integrate many other middleware technologies without undue effort. Figure 3 shows how SIML resolves the following challenges to generate a gateway given an existing CCM application:

Resolving challenge 1. Choosing an appropriate level of integration.

As mentioned in Section 2.1, determining the right level of integration requires expertise in all the different technologies being integrated. To allow interactions between CCM components and Web Services, SIML defines interactions between ports of CCM components and ports exposed by the Web Services. Since SIML also automates the generation of the glue code, some choices with respect to the level of integration, e.g., mapping of a CCM port to a Web Service port, are pre-determined, while other decisions, e.g., aggregation of more than one CCM component into a single Web Service, are customizable. SIML thus extends the list of valid interactions of both CCM com-

ponents and Web Services, which is an example of a composite DSML defining interactions that does not exist in its sub-DSMLs. SIML can also partition a large system into hierarchies via the concept of “modules,” which can be either CCM components (and assemblies of CCM components) or Web Services. SIML’s architecture can be enhanced to support integration of many middleware technologies, by extending the list of interactions defined by SIML to integrate new technologies. For example, SIML could be extended to support interactions between CCM and EJB, or even between Web Services and EJB. Extending SIML to support other technologies, *e.g.* EJB, requires specification of a DSML that describes the elements and interactions of EJB. Once the DSML for EJB is specified, it can be imported into SIML as a library while also assigning a new namespace to it; the creation of a new namespace prevents any clash between the type systems *e.g.*, between a CCM component and EJB component. Interactions between elements of CCM and EJB can then be defined in the composite DSML. From these new interactions, generative techniques (as explained in resolution to Challenge 3 below) can be applied to automate the task of integration.

Resolving challenge 2. Reconciling differences in interface specifications. To map interfaces between CCM and Web Services, SIML provides a tool called IDL2WSDL, which automatically converts any valid CORBA IDL file to a corresponding WSDL file. As part of this conversion process, IDL2WSDL performs both *datatype mapping*, which maps CORBA datatypes to WSDL datatypes, and *exception mapping*, which maps both CORBA exceptions to WSDL faults. System integrators are therefore relieved from the intricacies of the mapping. As shown in Figure 3, both IDL and WSDL can also be imported into the DSML environment corresponding to CCM (PICML) and Web Services (WSML), allowing integrators to define interactions between CCM components and Web Services. SIML also supports *language mapping* between ISO/ANSI C++ and Microsoft C++/CLI, which is the .NET framework extension to C++. SIML therefore automates much of the tedious and error-prone details of the interface mapping, thereby allowing system integrators to focus largely on the business logic of the application being integrated.

Resolving challenge 3. Managing differences in implementation technologies. While the rules defined in SIML allow definition of interaction at the modeling level, this feature is not very useful if these definitions cannot be translated into runtime entities that actually perform the interactions. SIML therefore generates *resource adapters*, which automatically convert SOAP requests into IIOP requests, and vice-versa. A *resource adapter* in SIML is implemented as a gateway. SIML allows system integrators to define connections between ports of a CCM component and a Web Service, as shown in Figure 3. These connec-

tions are then used by a *model interpreter*, which automatically determines the operation/method signatures of operations/methods of the ports on either end of a connection, and uses this information to automatically generate a gateway. The generated gateway contains all the “glue code” necessary to perform datatype mapping, exception mapping, and language mapping between CCM and Web Services. The gateway generator is configurable and can currently generate Web Service gateways for two different implementation of Web Services: GSOAP [24] and Microsoft ASP.NET. The generated gateway also performs the necessary *protocol mapping* (*i.e.*, between IIOP and SOAP) and *discovery mapping* (*i.e.* automatically connecting to a Naming Service to obtain object references to CCM components). Our initial implementation does not yet support *QoS mapping*, which is the focus of future work, as described in Section 6.

Resolving challenge 4. Managing deployment of subsystems. After the necessary integration gateways have been generated, system integrators also need to deploy and configure the application and the middleware using a variety of metadata in the form of XML descriptors. Since SIML is built using (meta)model composition it can automatically use the tools developed for the sub-DSMLs (*i.e.*, PICML to handle deployment of CCM applications and WSML to handle deployment of Web Services) directly from within SIML.

By encapsulating the required resource adapters inside a Web Service or CCM component, SIML allows reuse of deployment techniques available for any given middleware system. System integrators therefore do not need to deploy resource adapters separately. While this approach works for in-process resource adapters (such as those generated by SIML), out-of-process resource adapters need support from a deployment descriptor generator. Since SIML is a DSML itself, this support could be added to SIML so it can generate deployment support for out-of-process resource adapters.

Resolving challenge 5. Dealing with interoperability issues. Since knowledge of the underlying middleware technologies is built into SIML, it can automatically compensate for incompatibilities during design time. For example, IDL2WSDL allows generation of WSDL that supports SOAP RPC encoding or an interoperable subset defined in the WS-I Basic Profile. System integrators therefore are better prepared to handle incompatibilities that only show up during integration testing. SIML can also define constraints on WSDL definition as prescribed by the WS-I Basic Profile, so that violations can also be checked at modeling time. Similarly, gateway generation can add workarounds for quirks of particular implementations automatically, thereby relieving system integrators from finding these problems during final integration testing. The automation of gateway generation also scales the integration activity since develop-

Integration Activity	Supported?	Automated?	Level of Automation (# of distinct steps)					
			Using SIML			Using Native Tools		
			Design	Implementation	Tool Use	Design	Implementation	Tool Use
Integration Design	Yes	No	0	0	1	1	1	0
Interface Mapping								
DataType Mapping	Yes	Yes	0	0	1	1	1	0
Exception Mapping	Yes	Yes	0	0	1	1	1	0
Language Mapping	Yes	Yes	0	0	1	1	1	0
Technology Mapping								
Protocol Mapping	Yes	Yes	0	0	1	1	1	0
Discovery Mapping	Yes	Yes	0	0	1	1	1	0
QoS Mapping	No	No	1	1	0	1	1	0
Deployment Mapping								
Descriptor Generation	Yes	Yes	0	0	1	0	0	1
Gateway Placement	No	No	1	1	0	1	1	0
Interoperability Mapping	Yes	Yes	0	0	1	0	1	0

Table 1. Evaluating Functional Integration using SIML

ers need not write system specific integration code. In addition, SIML allows evolution of the integrated system by incrementally adding more components, or targeting different middleware implementations as future needs dictate.

4.3. Evaluating SIML

To evaluate the benefits of SIML, we first define a taxonomy for evaluating technologies that assist the functional integration of CCM and Web Services. We then use this taxonomy to compare SIML with tools that are supplied by vendors for either technology, referred to in Table 1 as *Native tools*. Examples of native tools include the Microsoft Visual Studio and the IBM Eclipse suite, which developers using middleware technologies like .NET and EJB are likely to use. This table depicts the different mapping activities described in Section 4.2 that are typical in functional integration of middleware systems. For each activity the table describes the level of support in SIML and whether the activity is automated. It also describes the level of automation measured as the number of distinct steps performed by a system integrator. Table 1 further decomposes the level of automation into three broad categories: (1) *design*, which denotes that system integrators need to perform a design activity that might include domain analysis, requirement analysis, etc., (2) *implementation*, which denotes that system integrators need to implement some functionality usually by writing code, and (3) *tool use*, which denotes that a tool needs to be used by the system integrators to perform that activity. This categorization assigns a weight commensurate to the skills of the individual responsible for carrying out the task in a particular organization.

Our taxonomy also assumes that design and implementation are orders of magnitude more difficult/time-consuming than tool use. In Table 1, therefore, multiple activities of the same category are considered equal, since the magnitude difference will likely dwarf any small number of steps

of any particular category. Thus the table uses 1 to indicate one or more, *i.e.*, $1 \dots n$ steps, and 0 to indicate that the effort is automated. To estimate the amount of effort required, we sum up each of the three columns (*i.e.*, design, implementation, and tool) and then multiply the result by the weight assigned to each category. For example, a reasonable assignment of weight for these activities might be 10, 5 and 1, for each of design, implementation and tool use. With this assignment, we can see that using SIML requires $2 \times 10 + 2 \times 5 + 8 \times 1 = 38$ distinct steps to achieve functional integration. In comparison, using just the native tools would result in $8 \times 10 + 9 \times 5 + 1 \times 1 = 126$ distinct steps to achieve the same. It should be noted that the number of steps will get reduced drastically as (and when) native tools add support for integration activities.

The numbers in Table 1 are for each unique unit of work per unique pair of source and target technologies, *i.e.*, for a single datatype mapping, a single exception mapping, a single protocol mapping. To calculate the total cost of integration, we must take into account both the number of distinct types/exceptions/languages, and the number of unique pairs of technologies being integrated.

Since SIML allows hierarchical composition of the integration infrastructure, the integration architecture scales along with the increase in the number of components. While the generative techniques applied to generate the gateways scale with the number of components in the system, when the number of components increases to thousands of components, the limitations of visual design tools tend to show up. To overcome the issues with scalability of modeling techniques, we have applied techniques like aspect-oriented weaving of domain-specific models [2] in prior efforts. Such techniques can be applied to automate the modeling activities in SIML in the presence of large number of components, since SIML itself is a domain-specific language for integration.

Table 1 shows that SIML helps reduce the effort by re-

ducing the design and/or implementation activities associated with integration to ordinary tool usage activities. For example, SIML effectively reduces the design and implementation effort required to perform the datatype, exception and language mapping, to a single step of tool use. This table also shows that similar gains can be achieved for complex tasks, such as protocol mapping (conversion between IIOP and SOAP in this case) and discovery mapping (conversion between CORBA Object References and Web Service URIs). Finally, the table reveals current gaps in our toolchain, *i.e.*, SIML does not perform QoS mapping or help with placement of resource adapters (or gateways), which remains as future work.

5. Related Work

This section surveys the technologies that provide the context of our work on system integration in the domain of large-scale distributed enterprise systems. We classify techniques and tools in the integration space according to the role played by the technique/tool in system integration.

Integration evaluation tools enable system integrators to specify the systems/technologies being integrated and evaluate the integration strategy and tools used to achieve integration. For example, IBM's WebSphere [9] supports modeling of integration activities and runs simulations of the data that is exchanged between the different participants to help predict the effects of the integration. While these tools help identify potential integration problems and evaluate the overall integration strategy, they do not replace the actual task of integration itself since these tools use simulation-/emulation-based abstractions of the actual systems. SIML's role is complementary to these tools: once the integration evaluation has been done using these tools, SIML can be used to design the integration, as well as generating the various artifacts required for integration.

Integration design tools. OMG's UML profile for Enterprise Application Integration (EAI) [18] defines a Meta Object Facility (MOF) [17] based metamodel for collaboration modeling, as well as activity modeling. MOF provides facilities for modeling the integration architecture focusing on connectivity, composition and behavior. The EAI UML profile also defines a MOF-based standardized data format to be used by the different systems to exchange data during integration, which is achieved by defining an EAI application metamodel that handles interfaces and metamodels for programming languages such as C, C++, PL/I and COBOL, to aid the automation of transformation. While standardizing on MOF is a step in the right direction, the lack of widespread support for MOF by various tools, and the differences between versions of XML Metadata Interchange (XMI) support in tools lead to problems in practice. Existing integration design tools provide limited support for in-

terface mapping by generating stubs and skeletons for facilitating interface mapping, and perform protocol mapping. However, key activities like discovery mapping, and deployment mapping still needs to be programmed by the system integrator. Thus the primary difference between SIML and these tools is that SIML not only allows such integration design, but also automates the generation of key integration artifacts, such as gateways, reducing the amount of effort required to develop and deploy the systems.

Integration patterns [23] provides guidance to system integrators in the form of best patterns and practices with examples of using a particular vendor's products. [8] catalogs common integration patterns with an emphasis on system integration via Message-Oriented Middleware (MOM) using different commercial products. These efforts do not directly provide tools for integration, but instead provide critical guidance to using existing tools to achieve integration. We are enhancing SIML to support modeling integration patterns and using them to enhance the generative capabilities of SIML to enable widely-accepted solutions to common integration problems.

Resource adapters are used during integration to transform data and services exposed by service producers to a form that is amenable to service consumers. While existing standards (such as the Java Messaging Specification [22] and J2EE Connector Architecture Specification [16]) and tools (such as IBM's MQSeries [10]) provide the architectural framework for performing the required adaptations, these tools approach the integration from a middleware and programming perspective, *i.e.*, system integrators are still required to handcraft the "glue" code that invokes the resource adapter frameworks to connect system components together. In contrast, SIML uses syntactic information present in the DSMLs to automatically perform the required mapping/adaptation by generating the necessary "glue" code, and relies on user input only for tool use.

Integration frameworks. Composition in the context of the semantic web and the Web Ontology Language (OWL) [5] has focused on composition of services from unambiguous, formal descriptions of capabilities as exposed by services on the web. Research on service composition has focused on automation and dynamism [19], optimizing the composition such that it is QoS-aware [26], as well as integration on large-scale "system-of-systems" like the GRID [7]. Since these automated composition techniques rely on unambiguous, formal representations of capabilities, system integrators need to make their legacy systems available as Web Services or provide alternate formal mappings of capabilities of the system to be integrated, which may not always be feasible. Our approach to (meta)model composition, however, is not restricted to a single domain, though the semantics are bound at design time. While both ap-

proaches rely on metadata, SIML's use of metadata focuses on the generative capabilities possible rather than on the semantic knowledge extracted from metadata.

Integration quality analysis. As the integration process evolves, it is necessary to validate whether the results are satisfactory from functional and QoS perspectives. Research on QoS issues associated with integration has yielded languages and infrastructure for evaluating Service-Level Agreements (SLAs). Examples include the Web Service Level Agreement language (WSLA) [14] framework, which defines an architecture to define service-level agreements using an XML Schema, and provides associated infrastructure to monitor the conformance of the running system to the desired SLA. Other efforts have focused on defining processes for distributed continuous quality assurance [25] of integrated systems to identify the impact on performance during system evolution. Information from these analysis tools should be incorporated into future integration activities. While these tools can be used to provide input to design-time integration activities, they themselves do not support automated feedback loops. We are adding support for modeling SLAs in SIML to allow evaluation of SLAs before/after integration.

6. Concluding Remarks

The development of enterprise distributed systems increasingly involves more integration of existing COTS software and less in-house development from scratch. With the increase in capabilities of COTS component middleware technologies, the complexity of integration of systems built upon such frameworks is also increasing. This paper shows how a model-driven approach to functional integration of component middleware technologies enhances conventional approaches to system integration, which are tedious, error-prone, and non-scalable for enterprise distributed systems. We then show how DSMLs and (meta)model composition can help to address these limitations.

To demonstrate the viability of our approach, we enhanced support for composition of DSMLs in GME. Using this new capability, we developed the *System Integration Modeling Language* (SIML), which is a DSML composed from two other DSMLs, the CCM profile of Platform-Independent Component Modeling Language (PICML) and the Web Services Modeling Language (WSML). Finally, we evaluated the benefits of our approach by generating a gateway from the model, which automates key steps needed to functionally integrate CCM components with Web Services.

The following is a summary of lessons learned thus far from our work applying (meta)model composition to inte-

grate heterogeneous middleware technologies:

- **Integration tools are becoming as essential as design tools.** SIML is designed to bridge the gap between existing *component technologies* (in which the majority of software systems are built) and *integration middleware* (which facilitate the integration of such systems). SIML elevates the activity of integration to the same level as system design by providing tools which allow integration design of systems built using heterogeneous middleware technologies. Since SIML is a DSML, it can potentially be used as the infrastructure to define constraints on the actual integration process itself, thereby allowing evaluation of service-level agreements prior to the actual integration itself.

- **Automating key portions of the integration process is critical to building large-scale distributed systems.** Compared with conventional approaches, our model-driven approach to system integration automates key aspects of system integration, including gateway “glue code” generation, metadata management, and design-time support for expressing unique domain and/or implementation assumptions. It supports seamless migration of existing investment in models and allows incremental integration of new systems. Moreover, our model-driven approach is general-purpose and can be applied to tool-chains other than GME, as well as help integrate systems other than CCM or Web Services.

- **QoS integration is a complex problem, and requires additional R&D advances.** Though SIML helped map functional aspects of a system from a source technology to a target technology, our work is not complete until the non-functional QoS-related aspects of a system also map seamlessly. For example, technologies like the Real-time CORBA Component Model (RT-CCM) [25] support many QoS-related features (such as thread pools, lanes, priority banded connections, and standard static/dynamic scheduling services) that allow system developers to configure the middleware to build systems with desired QoS features. When systems based on RT-CCM are integrated with other technologies, it is critical to automatically map the QoS-related features used by an application in the source technology to the set of QoS features available in the target technology. For example, a number of specifications have been released for Web Services that target QoS features, such as reliable messaging, security, and notification. The focus of our future efforts in integration involves extending SIML to automatically map QoS features from one technology to another using DSMLs, such that the integration is automated in all aspects – both functional and non-functional.

Instructions for downloading SIML and GME are available at www.dre.vanderbilt.edu/cosmic.

References

- [1] Ákos Lédeczi, G. Nordstrom, G. Karsai, P. Volgyesi, and M. Maroti. On Metamodel Composition. In *Proceedings of the 2001 IEEE International Conference on Control Applications (CCA)*, pages 756–760, Mexico City, Mexico, 2001. IEEE. 1.2, 3, 3
- [2] K. Balasubramanian, A. Gokhale, Y. Lin, J. Zhang, and J. Gray. Weaving Deployment Aspects into Domain-Specific Models. *International Journal on Software Engineering and Knowledge Engineering (IJSEKE)*, 16(3), June 2006. 4.3
- [3] K. Ballinger, D. Ehnebuske, C. Ferris, M. Gudgin, C. K. Liu, M. Nottingham, and P. Yendluri. WS-I Basic Profile. www.ws-i.org/Profiles/BasicProfile-1.1.html, April 2006. 2.1
- [4] C. Britton and P. Bye. *IT Architectures and Middleware: Strategies for Building Large, Integrated Systems*. Addison-Wesley Professional, May 2004. 1.1
- [5] W. W. W. Consortium. Web Ontology Language. www.w3.org/2004/OWL/, Feb 2004. 5
- [6] M. Emerson and J. Sztipanovits. Techniques for metamodel composition. In *The 6th OOPSLA Workshop on Domain-Specific Modeling, OOPSLA 2006*, Portland, OR, Oct 2006. ACM. 3
- [7] I. Foster, C. Kesselman, J. M. Nick, and S. Tuecke. Grid Services for Distributed System Integration. *Computer*, 35(6):37–46, 2002. 5
- [8] G. Hohpe and B. Woolf. *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*. Addison-Wesley Professional, October 2003. 1.1, 5
- [9] IBM. WebSphere. www.ibm.com/software/info1/websphere/index.jsp. 5
- [10] IBM. MQSeries Family. www-4.ibm.com/software/ts/mqseries/, 1999. 5
- [11] G. Karsai, S. Neema, B. Abbott, and D. Sharp. A Modeling Language and Its Supporting Tools for Avionics Systems. In *Proceedings of 21st Digital Avionics Systems Conf.*, Aug. 2002. 1.2
- [12] G. Karsai, J. Sztipanovits, A. Ledeczi, and T. Bapty. Model-Integrated Development of Embedded Software. *Proceedings of the IEEE*, 91(1):145–164, Jan. 2003. 1.2
- [13] A. Ledeczi, A. Bakay, M. Maroti, P. Volgysei, G. Nordstrom, J. Sprinkle, and G. Karsai. Composing Domain-Specific Design Environments. *IEEE Computer*, pages 44–51, November 2001. 1.2, 3
- [14] H. Ludwig, A. Keller, A. Dan, R. P. King, and R. Franck. Web Service Level Agreement Language Specification. researchweb.watson.ibm.com/wsla/documents.html, January 2003. 5
- [15] B. Meyer. Applying Design By Contract. *Computer (IEEE)*, 25(10):40–51, Oct. 1992. 3
- [16] S. Microsystems. J2EE Connector Architecture Specification. java.sun.com/j2ee/connector/, November 2003. 5
- [17] Object Management Group. *MetaObject Facility (MOF) 2.0 Core Specification*, OMG Document ptc/03-10-04 edition, Oct. 2003. 5
- [18] Object Management Group. *UML Profile for Enterprise Application Integration (EAI)*, omg document formal/04-03-26 edition, March 2004. 5
- [19] S. R. Ponnokanti and A. Fox. SWORD: A Developer Toolkit for Web Service Composition. Jan. 01 2002. 5
- [20] D. C. Schmidt. Model-Driven Engineering. *IEEE Computer*, 39(2):41–47, 2006. 1.2
- [21] J. A. Stankovic, H. Wang, M. Humphrey, R. Zhu, R. Pooralingam, and C. Lu. VEST: Virginia Embedded Systems Toolkit. In *Proceedings of the IEEE Real-time Embedded Systems Workshop*, London, UK, Dec. 2001. IEEE. 1.2
- [22] SUN. Java Messaging Service Specification. java.sun.com/products/jms/, 2002. 5
- [23] D. Trowbridge, U. Roxburgh, G. Hohpe, D. Manolescu, and E. G. Nadhan. Integration Patterns. msdn.microsoft.com/library/default.asp?url=/library/en-us/dnpag/html/intpatt.asp, June 2004. 5
- [24] R. van Engelen and K. Gallivan. The gSOAP Toolkit for Web Services and Peer-to-Peer Computing Networks. In *CC-GRID*, pages 128–135. IEEE Computer Society, 2002. 4.2
- [25] N. Wang and C. Gill. Improving Real-time System Configuration via a QoS-aware CORBA Component Model. In *Hawaii International Conference on System Sciences, Software Technology Track, Distributed Object and Component-based Software Systems Minitrack, HICSS 2004*, Kona, HI, Jan. 2004. HICSS. 5, 6
- [26] L. Zeng, B. Benatallah, A. H. Ngu, M. Dumas, J. Kalagnanam, and H. Chang. QoS-Aware Middleware for Web Services Composition. *IEEE Trans. Softw. Eng.*, 30(5):311–327, 2004. 5