**Noname manuscript No.**
(will be inserted by the editor)

# A Model Transformations-based Approach to Automating Middleware QoS Configurations

**Amogh Kavimandan**[1⋆], **Aniruddha Gokhale**[1⋆⋆], **Anantha Narayanan**[1], **Gabor Karsai**[1]

Dept. of EECS, Vanderbilt University, Nashville, TN 37235.

Received: Feb 1, 2009

**Abstract** The development and operational life-cycles of distributed real-time and embedded (DRE) systems can be improved using component middleware due to its support for rapid assembly and deployment of large applications. The flexibility of component middleware can, however, complicate DRE system development since non functional system properties, such as system quality of service (QoS), depends on the effective configuration of the middleware. DRE system developers often lack detailed knowledge of the underlying middleware, which makes it challenging for them to map the domain-specific QoS requirements into the right set of QoS configurations of the middleware. This paper describes a technique based on model transformations that addresses these challenges. Our technique enables developers to use domain-specific abstractions to specify their QoS requirements, which are transformed into middleware QoS configurations using model transformation algorithms we developed. Verifying the correctness of the transformation process itself and that of the generated QoS configurations is addressed using structural correspondence and automated model checking, respectively. A proof-of-concept validation of the effectiveness of our technique to meet the end-to-end QoS requirements is presented in the context of a representative DRE system.

**Keywords:** Model transformation, DSMLs, QoS configurations, middleware.

## 1 Introduction

Contemporary component middleware technologies, such as Enterprise Java Beans (EJB) and CORBA Component Model (CCM), help to decouple the development of application logic from the provisioning of non functional properties of the system, such as quality of service (QoS) properties, which are key to the correct functioning of distributed real-time and embedded (DRE) systems. Examples of DRE system QoS properties that can be configured in the middleware includes different types of concurrency, handling multiple levels of priorities of system tasks, publish/subscribe event-driven communication mechanisms, reliability, security, and multiple scheduling algorithms, among others.

Although component middleware separates the QoS configuration complexity from the application logic, the need to support multiple kinds of QoS properties has made the middleware itself very complex, which in turn makes the middleware QoS configuration activity very complex. Middleware QoS configuration involves mapping the system-level *QoS policies*—which are dictated by domain requirements—onto the solution space comprising the *QoS mechanisms* for tuning the underlying middleware.

Examples of domain-level QoS requirements include (1) the degree of concurrency required to provide a service, (2) the priorities at which the different components should run, (3) the alternate protocols that can be used to request a service, (4) the granularity of sharing among the application components of the underlying resources, such as transport level connections, (5) the number and size of outstanding requests that are permissible at any instant in time, and (6) the maximum and minimum amount of time to wait for completion of requests, among many others. All these domain-specific QoS requirements must eventually be mapped onto the right middleware QoS configurations, which includes choosing the right set of middleware configuration parameters and assigning the right values to them.

An additional dimension of complexity stems from the need to perform QoS configurations at different time scales. This includes *statically*, *e.g.*, directly hard coded into the application or middleware; *semi-statically*, *e.g.*, configured at deployment time using metadata descriptors; and/or *dynamically*, *e.g.*, by modifying QoS configurations at runtime.

The complexity involved in the middleware QoS configuration process raises the following questions:

- How can the domain-specific QoS requirements of the system be mapped onto QoS configurations of the under-

⋆ Author has since graduated from Vanderbilt University
⋆⋆ Contact Author: a.gokhale@vanderbilt.edu

lying middleware particularly by DRE system developers who are domain experts but seldom possess detailed knowledge of the middleware and its configurability?

- How are the right set of configuration parameters determined and how are valid values for the selected set of QoS configuration parameters determined? Are there any *patterns of usage i.e.*, best practices in configuration that can be used across a variety of application domains?
- How can dependency relationships between configuration parameters be resolved while ensuring that their interactions do not adversely impact QoS? These issues arise at individual component level (local) as well as at aggregate intermediate levels, such as component assemblies, all the way through the entire application (global) and at different time scales, *i.e.*, design-, deployment- and run-time.

Without scientific techniques and effective tools to address these questions, existing approaches to middleware QoS configurations will result in QoS mis-configurations that are hard to analyze and debug. As a result, failures will stem from a new class of configuration errors rather than (just) traditional design/implementation errors or resource failures.

**Solution approach → Model-to-model Transformations:** We address these questions by presenting a technique based on model-to-model transformations, which is realized in the context of a tool-chain called *QUality of service pICKER* (QUICKER) [10]. We present the scientific principles behind the model-to-model transformations in QUICKER concretely in the context of two domain-specific modeling languages (DSMLs).

The first language is called *Platform-Independent Component Modeling Language* (PICML) [3]. PICML enables developers of component-based DRE systems to annotate component-based application models developed in PICML with QoS policies. These policies are specified at a higher-level of abstraction using *platform-independent* models, rather than using low-level, platform-specific configuration options typically found in middleware configuration files. QUICKER thus allows flexibility in mapping the same QoS policy to other middleware technologies.

QUICKER uses model-to-model transformation techniques [5] to translate the platform-independent PICML specifications of QoS policies into middleware QoS configurations modeled in another DSML called the *Component QoS Modeling Language* (CQML). CQML enables the modeling of platform-specific QoS configurations provided by the CIAO [6] real-time CORBA Component Model middleware.

QUICKER makes the following contributions in addressing the questions raised above:

- QUICKER provides DSMLs that enable DRE system developers to use intuitive abstractions to model system QoS policies. The model-to-model transformations subsequently automate the middleware QoS configuration process.
- The model-to-model transformations in QUICKER determine the right set of configuration parameters and their values. We use structural correspondence to verify the correctness of our transformation process.
- Dependencies among the configuration options and the correctness of the generated configurations at all time scales is resolved using automated model checking provided by the Bogor [23] model-checking framework.

The remainder of this paper is organized as follows: Section 2 describes the key challenges in QoS configuration of component middleware; Section 3 describes the QUICKER model-to-model transformation approach; Section 4 discusses how we have employed structural correspondence and model-checking techniques to verify the correctness of the QoS configuration process; Section 5 provides empirical validation of the generated QoS configurations for a representative DRE system; Section 6 compares QUICKER with related work on model-driven QoS config/-uration/adapta/-tion; and Section 7 presents concluding remarks.

## 2 Sources of Complexity in Configuring Middleware for DRE Systems

This section describes the challenges in configuring middleware, which is needed to meet the QoS requirements of DRE systems. To better explain these challenges, we describe them in the context of a representative middleware – in our case real-time CORBA [20] and the CIAO CORBA component middleware that leverages it.

### *2.1 Overview of Real-time CORBA and Component Middleware*

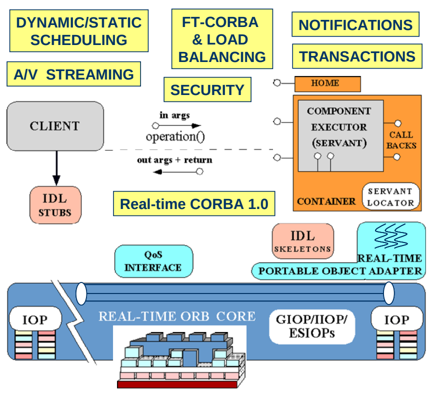


Fig. 1: **Real-time CORBA Middleware**

Figure 1 illustrates the Real-time CORBA (RTCORBA) middleware architecture, which is leveraged by the CIAO component middleware platform. RTCORBA extends traditional CORBA artifacts, such as (a) the object request broker (ORB), which mediates the request handling between clients and servers, (b) the portable object adapter (POA), which manages the lifecycle of CORBA objects, (c) stubs and skeletons, which are generated by an interface definition language (IDL) compiler that hide the distribution aspects from the communicating entities, with real-time policies and interfaces.

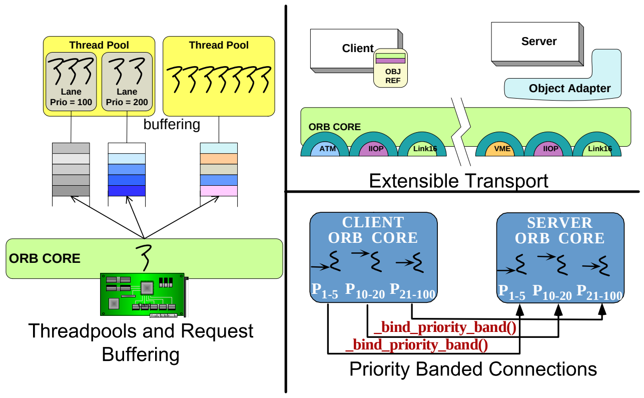


Fig. 2: **Real-time CORBA Middleware Features**

RTCORBA defines standard interfaces and QoS policies shown in Figure 2 that allow applications to configure and control (1) *processor resources* via thread pools, priority mechanisms, intra-process mutexes, and a global scheduling service, (2) *communication resources* via protocol properties and explicit bindings, and (3) *memory resources* via buffering requests in queues and bounding the size of thread pools. Applications typically specify these real-time QoS policies along with other policies when they call standard CORBA operations, such as `create_POA` or `validate_connection`. For example, the priority at which requests must be handled can be propagated from the client to the server (the `CLIENT_PROPAGATED` model) or declared by the server (the `SERVER_DECLARED` model).

The CIAO component middleware is an implementation of the lightweight CORBA Component Model (LwCCM) [19] and leverages RTCORBA. DRE system developers can realize large-scale DRE systems by assembling and deploying CIAO-based LwCCM components. The applications within these DRE systems can use publish/subscribe communication semantics (by using the component event source and sink ports) or request/response communication semantics (by using the facet and receptacle ports). The QoS properties of these applications are realized only by appropriately configuring the CIAO middleware and its RTCORBA policies.

### *2.2 Middleware Configuration Challenges*

We now present the sources of complexities in configuring the middleware in accordance to the domain-specified QoS requirements of the DRE system.

**Challenge 1. Inherent complexity in translating QoS policies to QoS configuration options.** Translating QoS policies into QoS configuration options is hard because it must transform semantics from the application domain to the semantics of the underlying component middleware. QoS-enabled component middleware like CIAO leverage RTCORBA mechanisms to configure (1) *processor resources*, such as portable priorities, end-to-end priority propagation, thread pools, distributable threads and schedulers, (2) *communication resources*, such as protocol properties and explicit binding of connections, and (3) *memory resources*, such as buffering of requests. To translate the domain-specified QoS policies into QoS mechanisms by configuring the QoS options, system developers need a thorough understanding of the underlying middleware platforms. For example, they must be able to determine the right model to handle priorities, and determine the right concurrency model.

**Challenge 2. Satisfying pre/post conditions and invariants in QoS configuration.** Assuming that a domain expert can translate the QoS policies into a subset of QoS configuration options, it is also necessary to understand the pre-conditions, invariants, and post-conditions of the different QoS configuration options since they affect middleware behavior. For example, to create `thread pool-with-lanes` in a server component, the following (non-exhaustive) pre-/post-conditions and invariants must be satisfied:

- *Pre-conditions.* A real-time POA created within a real-time ORB must be available, and the range of priorities (for the different lanes) and the type of priority mapping scheme chosen must be compatible, *i.e.*, within the limits.
- *Post-conditions.* A thread pool-with-lanes corresponding to the different priorities, along with the requested number of static (pre-defined) threads must be made available for use.
- *Invariants.* The real-time ORB should be able to match incoming request priorities to the corresponding lanes, and must always handle incoming requests for higher priority lanes before incoming requests for lower priority lanes.

This problem is exacerbated by the plethora of options and choices of valid values for each option, as well as by the fact that choosing one value for a particular option may have side effects on other options. These side effects are sometimes manifested as overt failures, such as failure to perform a mapping of CORBA priority to the underlying OS priority because of insufficient priorities in the OS to support the choice of priority mapping scheme, *e.g.*, *direct* mapping. They may also be manifested, however, as hard-to-reproduce and/or debug runtime failures that only emerge during field testing, or after deployment, which are much harder to detect and fix.

**Challenge 3. Resolving dependencies between QoS configuration options.** Even with a thorough understanding of middleware QoS configuration options, manual configuration of QoS policies cannot scale as the number of entities to configure increases. This lack of scalability stems from dependencies between the different QoS configuration options of each component, such as the dependency between the CORBA priority of a component, the chosen priority mapping scheme (to map CORBA priority to native OS priority), and the priority-banded connections policy (which selects the appropriate connection to route requests based on the request invocation priority). As the number of components increases, the number of inter-component dependencies increases proportionally. If the components are connected, the side effect of the connection between components may also induce an inter-component option dependency. Since these dependencies can grow quadratically, it is infeasible for developers to manage these dependencies manually.

In DRE systems with many components, the effects of changing a QoS configuration option on a component may affect many other directly connected components, their connected neighbors and so on. These dependencies can rapidly degenerate into a very large number of QoS configurations. Keeping track of dependencies between options and propagating the changes in one option to all options affected by that change is critical during the QoS configuration phase.

Validating the values of the different QoS configuration options in isolation and together with connected components is critical to the successful deployment and ultimately the operation of DRE systems. Once again, it is hard to validate these values without automated tool support. Depending on the frequency of changes made to the domain-specified QoS requirements during the development process, empirically validating a change in QoS configuration options becomes time consuming at this scale, which slows down the design process considerably and permits subtle and pernicious errors to occur. What is needed is an automated tool support that can assure an application's evolution throughout its entire lifecycle.

**Challenge 4. Ensuring validity of QoS configuration options with changes in QoS policies.** QoS configuration options effect the non-functional behavior of a system, and thus are affected by changes in the system environment. For a DRE system to operate effectively in hostile environments, component middleware and their associated QoS configuration options may need to adapt to their current conditions. While it is useful to change QoS configuration options at runtime to effect changes in behavior (such as re-prioritizing or increasing/decreasing resource usage), such dynamic reconfigurations may incur another set of challenges.

In particular, not only must we handle static QoS configuration problems (such as checking validity of values and keeping track of dependencies), there is typically little leeway to accommodate misconfiguration at runtime. It is non-trivial to change a running system because the system might crash during reconfiguration due to misconfiguration of QoS options. Moreover, the reconfiguration process itself must be predictable for the reconfiguration to have the desired effect on system behavior. In a DRE system, for instance, a reconfiguration done too late may be worse than not performing a reconfiguration at all.

An exhaustive evaluation of possible choices of QoS configuration options and validation of the reconfigured state is too time consuming to perform at runtime and can delay the reconfiguration process itself, rendering it useless. Once again, tools are needed to help validate and automate this reconfiguration process.

The remainder of this paper shows how QUICKER helps to address these challenges.

## 3 The QUICKER QoS Configuration Process

Figure 3 shows the middleware QoS configuration approach adopted by QUICKER. DRE system developers use the *Requirements* domain-specific modeling language (DSML)/metamodel to specify the system QoS requirements. A specification of system QoS requirements modeled in this language acts as the source model for the model-to-model transformation algorithms. Middleware-specific QoS configuration options are generated as models of the *QoS Configurations* DSML, which serves as the target DSML in the transformation process.

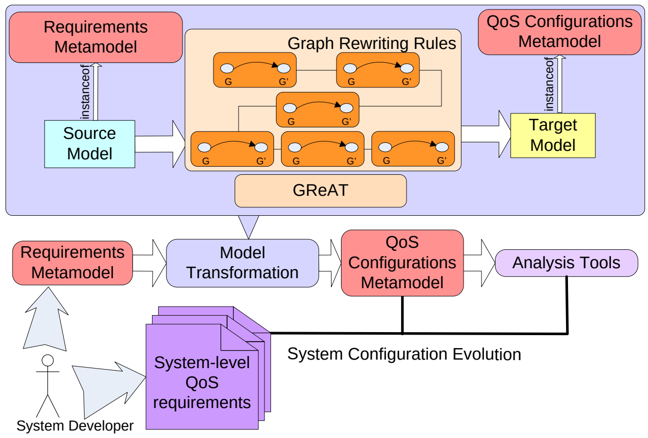


Fig. 3: **Model-driven QoS configuration process**

To demonstrate the viability of the QUICKER approach, we have used the Generic Modeling Environment (GME) [1] framework for developing the source and target DSMLS. GME provides a graphical user interface that can be used to define both the language semantics and system models that conform to the languages defined in it. The model-to-model transformations have been developed [11] using the Graph Rewriting And Transformation (GReAT) [9] framework. GReAT, which is implemented within the framework of GME, can be used to define transformation rules using its visual language, and executing these transformation rules for generating target models using the GReAT execution engine (GR-Engine).

### *3.1 Modeling QoS Requirements: The Source DSML*

QUICKER defines a `Requirement` element as a generalization of QoS requirements. As shown in Figure 4, source elements `Component`, `ComponentAssembly` or `Port` connections can be associated with a `Requirement` element. The modeling abstractions in QUICKER allow association of multiple source elements with the same `Requirement` as long as those source elements are of the same *type*. Moreover a `ComponentAssembly`'s `Requirement` is also associated with all the components contained in that `ComponentAssembly`. Such associations provide significant benefits in terms of flexibility in the creation of QoS requirements models and scalability of the models. The metamodels we describe below have been integrated with PICML using these associations; thus a single model of DRE system captures its entire QoS requirements specification.

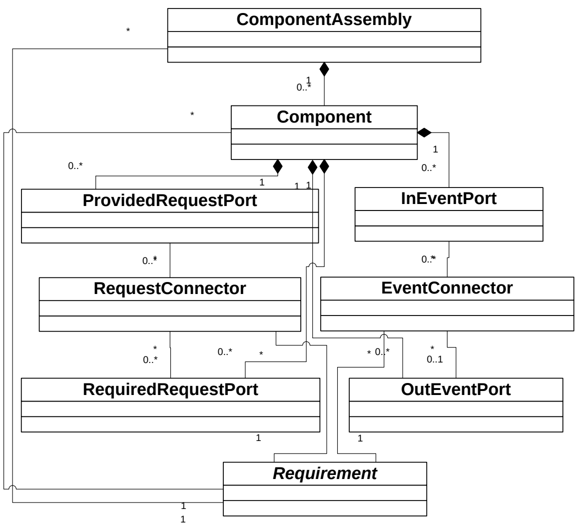


Fig. 4: **Simplified UML notation of QoS requirements associations in QUICKER**

Next we discuss the requirements specification across the following two real-time QoS dimensions: (1) RT request/response that is used to specify requirements for components and the synchronous connections between components, and (2) RT publish/subscribe service that is used to specify requirements for asynchronous event-based connections between components.

**Request/Response QoS requirements.** The request/response requirements have component-level granularity. A `RTRequirement` element, which is derived from `Requirement`, captures real-time requirements of a component and may have the following two attributes: (1) `fixed_priority_service_execution`, a server component Boolean property for specifying whether or not it modifies client service invocation priorities, and (2) `bursty_client_requests`, a server component Boolean property for specifying the profile of service invocations made by its client components. Figure 5 illustrates the relevant concrete syntax for modeling the requirements for the request/response real-time communication.

**Publish/subscribe QoS requirements.** We have modeled the requirements for real-time publish/subscribe event service to enable specification of QoS across asynchronous and anonymous interactions in component-based DRE systems. In the context of a publish/subscribe service, a **Subscriber** component subscribes to receive events from a **Publisher** component that generates events. Publisher (subscriber) component connects to a mediator entity, an **Event Channel**, to publish (subscribe to) events. Figure 6 illustrates the relevant concrete syntax for modeling the requirements of the event channel.

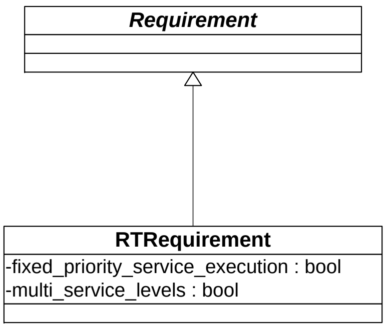


Fig. 5: **Simplified UML notation for Requirements of Real-time Request/Response**

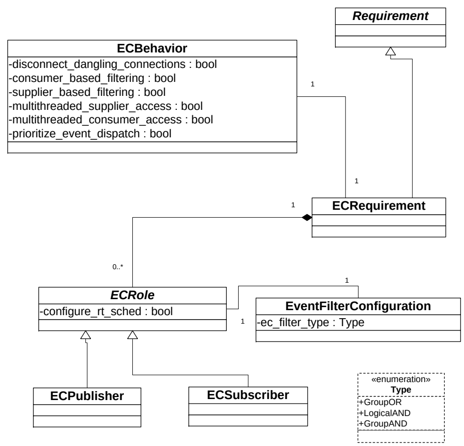


Fig. 6: **Simplified UML notation for Requirements of Event Channels**

A `ECRequirement` element is derived from `Requirement`. It models the properties of the event channel and can be used to specify the following QoS requirements: (1) `network_quality`, a connection-level property that captures the quality value of network used for running the application; (2) `connection_frequency`, a component-level property specifying the frequency at which the component (dis)connects with the publish/subscribe connection; (3) `event_distribution_ratio`, a connection-level property that specifies the ratio: $\frac{E^a{}_c}{E^s{}_c}$, where $E^a{}_c$ denotes number of events available for subscription at connection $c$ and $E^s{}_c$ denotes average number of events subscribed to at connection $c$ by each subscriber component.

These modeling capabilities are at a sufficiently high level of abstraction and are intuitive to be applied to a variety of publish/subscribe mechanisms. All the requirements have an enumerated data type with values `LO` and `HI`.

### *3.2 Modeling Middleware QoS Options: The Target DSML*

Rather than directly transforming source models of DRE system QoS policies into configuration descriptors in XML required for deploying it on the middleware, we chose to generate models of middleware-specific QoS options from these source models such that they can be used for further analysis such as model-checking QoS properties of the DRE system. This represents the Component QoS Modeling Language (C-QML). We have developed interpreters for parsing CQML system models and generating XML deployment descriptors in preparation for deploying the DRE system on the target environment.

**Request/Response QoS options.** A simplified UML notation of CQML modeling elements for Request/Response QoS options is shown in Figure 7.

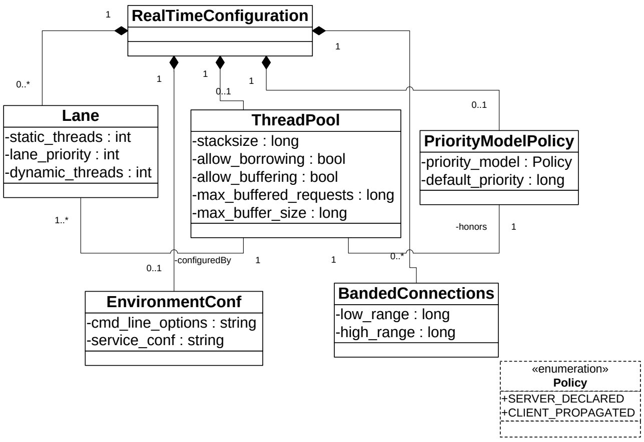


Fig. 7: **Simplified UML notation for RT-CCM Request/Response Configuration Options**

As shown, CQML defines the following elements corresponding to several RT-CCM configuration mechanisms: (1) `Lane`, which is a logical set of threads each one of which runs at `lane_priority` priority level. It is possible to configure *static* thread (*i.e.*, those that remain active till the system is running and *dynamic* thread (*i.e.*, those threads that are created and destroyed as required) numbers using `Lane` element; (2) `ThreadPool`, which controls various settings of `Lane` elements, or a group thereof. These settings include `stacksize` of threads, whether borrowing of threads across two `Lane` elements is allowed, and maximum resources assigned to buffer requests that cannot be immediately serviced; (3) `PriorityModelPolicy`, which controls the policy model that a particular `ThreadPool` follows. It can be set to either CLIENT_PROPAGATED if the invocation priority is preserved, or SERVER_DECLARED if the server component changes the priority of invocation; and (4) `BandedConnections`, which defines separate connections for individual (client) service invocations.

**Publish/subscribe QoS options.** For QoS configuration of asynchronous event communications, CQML defines the following elements: (1) `Publisher` and `Subscriber` modeling elements contain all the event source and sink settings, respectively. These include, for example, thread locks management mechanisms for publishers (subscribers) that are accessed by multi-threaded systems, and types of event filtering used, (2) `RTECFactory` element contains configurations specific to the event channel itself. These include, for example, event dispatching method that controls how events from publishers are forwarded to the respective subscribers, scheduling of events for delivery and other scheduler-related coordination, and handling of timeout events in order to forward them to respective subscribers, and (3) `FilterGroup` element that specifies strategies to group more than one filters together for publishers (subscribers).

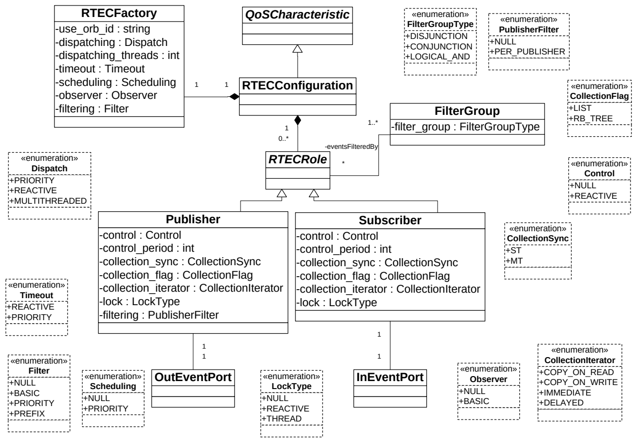


Fig. 8: **Simplified UML notation for RT-CCM Publish/-Subscribe Configuration Options**

Having a DSML such as CQML has the following advantages:

- The (generated) CQML models can be used for further analysis such as, for example, model-checking QoS properties of the DRE system.
- Using models of application at each step of software development increases traceability as QoS requirements are translated to low-level QoS options *i.e.*, to middleware descriptors required for deploying the DRE system.
- The CQML models are *closed* in terms of modification by system developers, and can only be (re-)generated by applying QUICKER model transformations on system QoS requirements models. Such a design choice simplifies software development and is a crucial step towards addressing productivity problem [13] at the middleware level.

### 3.3 Model-to-model Transformation Algorithms

The QUICKER model transformation rules have been defined in GReAT and are based on our past experiences in configuring QoS for component-based DRE systems. They are applicable to any system model that conforms to the Requirements DSML, and thus can be used by the system developers repetitively during the development and/or maintenance phase(s) of the DRE system. QUICKER model transformations preserve the granularity specified in the source models.

**Mapping real-time QoS requirements.** Let ${R_p}^o$ and ${R_p}^i$ denote, respectively, the set of outgoing (required/event source) and incoming (provided/event sink) ports of component $p \in P$. Let $S$ and $C$ be the sets of server and client components respectively and are given by:

$p \in S \; if \; {R_p}^i \neq \emptyset$ and $p \in C \; if \; {R_p}^o \neq \emptyset$

Algorithm 1 describes (non-exhaustive) RT-CCM QoS mappings in QUICKER. Lines 5-13 show the thread resource allocation scheme for server components. For every incoming port of a server component, the number of interface operations and client components are counted (lines 9 and 10). These counts are used by the auxillary function $ThreadResources$ to calculate the total threads required for handling all client service invocations at that server.

---

**Algorithm 1**: **Real-time QoS requirements mapping**

---

**Input**: set of client components $C$, set of server components $S$, set of bursty client components $B$, set of threadPool lanes $TPLanes$

1 **begin**
2 InterfaceOperationsCount $ioc$; ClientsCount $cc$;
3 IncomingPort $ip$; OutgoingPort $op$;ThreadCount $tc$;
4 Component $c$;set of Components $CPS$; Buffering $bf$;
5 **foreach** $p \in S$ **do**
6 $ioc \leftarrow 0$; $cc \leftarrow 0$; $tc \leftarrow 0$; $bf \leftarrow false$;
7 $CPS \leftarrow ClientComponents(p)$;
8 **foreach** $ip \in {R_p}^i$ **do**
9 $ioc \leftarrow ioc + countOperations(p, ip)$;
10 $cc \leftarrow cc + countClientComponents(p, ip)$;
11 **end**
12 $tc \leftarrow ThreadResources(ioc, cc)$;
13 $createTPLanes(p, tc)$;
14 **foreach** $c \in CPS$ **do**
15 **if** $c \in B$ **then**
16 $bf \leftarrow true$;
17 $assignThreadResources($
18 $TPLanes_p, c, tc)$;
19 $assignTPoolAttributes(TPLanes_p, bf)$;
20 $ioc \leftarrow 0$;
21 **foreach** $op \in {R_c}^o$ **do**
22 $ioc \leftarrow ioc + countOperations(c, op)$;
23 **end**
24 $createBands(c, ioc)$; $matchPriorities(p, c)$;
25 **end**
26 **end**
27 **end**

---

For handling bursts of client requests, server components should configure their thread pool to grow dynamically such that threads are created only when required.$assignThreadResources$ function is used to adjust the ratio of static and dynamic threads for a server, depending on whether its `bursty_client_requests` property is set to TRUE. In addition, lane borrowing feature at the server is set to TRUE such that the thread pool lanes across various priority levels can be borrowed. Finally, *PriorityBands* are configured and the their priority values are matched with server-side lane values in line 24.

**Mapping publish/subscribe QoS requirements.** Let ${PC_c}^s$ denote the synchronization mechanism, ${PC_c}^t$ denote the type, ${PC_c}^i$ denote the iterator in proxy collection $PC$ for component $c$, respectively. Let $L_c$ denote the locking policy, $CP_c$ denote control policy, $SF_c$ denote supplier-based filtering at component $c$, respectively. Algorithm 2 gives the (non-exhaustive) publish/subscribe QoS mappings.

---

**Algorithm 2**: **Publish/Subscribe service QoS requirements mapping**

---

**Input**: set of components $CPS$

1 **begin**
2 Component $c$; ThreadPoolLaneCount $lc$;
3 NetworkQuality $nq$;
4 **foreach** $c \in CPS$ **do**
5 $lc = countThreadResources(c)$;
$cf = connectionFrequency(c)$;
$nq = networkQuality(c)$;
$dr = eventDistributionRatio(c)$;
6 **if** $lc \neq 1$ **then**
7 ${PC_c}^s = MT$; $L_c = THREAD$;
8 **else**
9 ${PC_c}^s = ST$; $L_c = NULL$;
10 **end**
11 **if** $cf \neq LO$ **then**
12 ${PC_c}^t = LIST$; ${PC_c}^i = COPY\_ON\_READ$;
13 **else**
14 ${PC_c}^t = RB\_TREE$; ${PC_c}^s = COPY\_ON\_WRITE$;
15 **end**
16 **if** $nq \neq LO$ **then**
17 $CP_c = NULL$;
18 **else**
19 $CP_c = REACTIVE$;
20 **end**
21 **if** $c \in S$ **then**
22 **if** $dr \neq LO$ **then**
23 $SF_c = PER\_SUPPLIER$;
24 **else**
25 $SF_c = NULL$;
26 **end**
27 **end**
28 **end**

---

A publish/subscribe service has several settings for configuring the way collections of publisher and subscriber object references are created and accessed, which must be chosen appropriately for individual applications. Lines 6-9 in Algorithm 2 show how the choice of serialization mechanism is affected by the number of thread resources configured at component *c*.

The choice of the *type* of collection is based on the following: (1) RB_TREE data structure exhibits faster ($O(log(n))$) insertion and removal operations. Therefore, it is more suited for connections whose components have a high (dis)connection rate; (2) LIST data structure on the other hand, should be chosen in cases where iteration is frequent (and therefore, more crucial for efficient application execution) than modifications to it.

Lines 11-14 give the steps in algorithm that configure the collection type. Finally, REACTIVE policy is chosen for applications that use low-quality value network on Lines 16-19, which ensures that (publisher/subscriber) components are periodically polled for determining their states (*i.e.*, whether or not they are connected to the event channel).

### *3.4 Resolution of Challenges 1 & 3*

Target typed graph elements (*i.e.*, QoS options), are well-understood by middleware implementation experts. We expect that the QUICKER transformation algorithms 1 and 2 will be designed in terms of source and target typed graphs by these experts. DRE system developers will only need to think of their requirements at levels that are intuitive. Hence, they can describe their system QoS requirements using the modeling capabilities discussed in Section 3.1.

By providing platform-independent modeling elements in QUICKER and defining representational semantics that closely follow those of the system requirements, QUICKER allows system developers to describe system QoS using simple, intuitive notations. Further, model transformations defined in QUICKER automatically identify and deduce QoS configurations that are best suited to achieve the desired QoS for DRE systems being configured. This resolves Challenge 1 described in Section 2.2.

Challenge 3 described in Section 2.2 is partly resolved as follows. QUICKER transformation rules contain information about the semantics of the QoS options, their inter-dependencies, and how they affect the high-level QoS requirements of a DRE system and therefore are used to assign values to the subset of options chosen earlier. Further, QoS options semantics are known precisely during transformations, and thus QUICKER ensures preservation of the target typed graph semantics. Component interactions defined in input typed graph instance (*i.e.*, source model), along with the user-specified QoS requirements captured in that instance are used to completely generate an instance of the output graph.

## 4 Correctness of the Model Transformation Process

In any transformation process, verifying the correctness of the process itself and the generated artifacts are important. Verifying the correctness of our model-to-model transformation-based QoS configuration process entails verification of the correctness properties across the following two dimensions [12]: (1) correctness of the QoS mapping process, *i.e.*, providing assurance that the QoS configurations generated are equivalent to the QoS requirements from which these options are mapped. In our case, this translates to verification of the QoS transformations used, and (2) correctness of the generated QoS options themselves, *i.e.*, whether the configuration parameters themselves are correct and whether the individual values of these options are appropriate locally (*e.g.*, for a component) as well as globally (*e.g.*, for all dependent components). This section discusses the correctness of our transformation process along the above two dimensions.

### *4.1 Verifying the Correctness of the Model Transformations*

To provide an assurance that the QoS requirements specifications are correctly mapped into the QoS configuration model, we have used a transformation verification technique based on structural correspondence described in [17]. Figure 9 shows an overview of the verification technique.

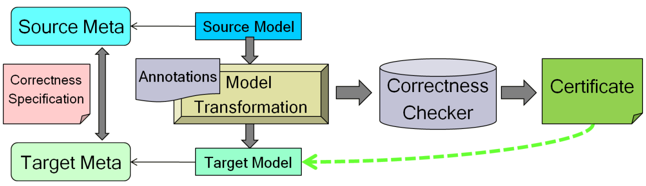


Fig. 9: **Verifying model transformations**

The source and target portions of the transformation are treated as typed, attributed graphs, called the source and target meta-models. The correctness of the transformation is specified as a relation between these graphs. Such a relation, called a structural correspondence, is specified by identifying pivot nodes in the meta-models and specifying what constitutes a correct transformation for these nodes. The transformation rules are extended to generate annotations along with the instance models. Once this setup has been completed, every execution of the transformation is followed by an automatic correctness checker, which uses the generated annotations to check if the correctness conditions were satisfied for the generated output model. If this check is satisfied, we can say that the output model instance is 'certified' correct.

Using structural correspondence, the verification consists of two phases: the specification of the correctness conditions, and the evaluation of the correctness. In the first phase, we identify important points in the transformation, and specify the structural correspondence rules for these points. From these rules, a model traverser is automatically generated, which will traverse and evaluate the correspondence rules on the instance models. This step needs to be performed only once.

The second phase involves invoking the model traverser after each execution of the model transformation. In this phase, the model instance being transformed is traversed, and the structural correspondence rules are evaluated at each relevant node. If any of the rules are not satisfied, it indicates that the model has not been transformed satisfactorily.

Structural correspondence rules are described using (1) specification of the correspondence condition itself, and (2) the rule path expressions, which are similar to XPath queries. Figure 10 shows how we have used cross-links in GReAT as means of specifying the correspondence condition between input and output language objects such that their equivalence can later be established.

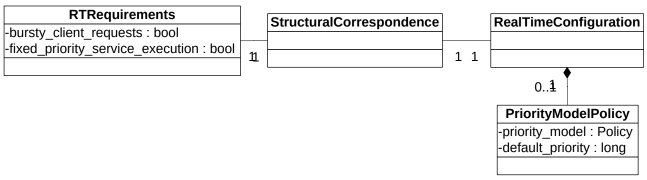


Fig. 10: **Structural correspondence using cross-links**

`RTRequirements` is an input language object that denotes real-time requirement specification for a component. It has a correspondence relation with `RealTimeConfiguration` output language object, indicated by presence of a cross-link between them in Figure 10. Additionally, one of the transformation rules in our QoS mapping algorithms states that if the Boolean attribute `fixed_priority_service_execution` of `RTRequirements` is set to TRUE in the input model, then `priority_model` attribute of `PriorityModelPolicy` object be set to SERVER_DECLARED in the output model. Otherwise `priority_model` should be set to CLIENT_PROPAGATED.

Moreover, if `priority_model` is set to SERVER_DECLARED for a component, `Lane` values at that component and `BandedConnection` values at its clients must match. In order to complete the correspondence rule specification, the above is encoded as a rule path expression as follows:

```
(RTRequirement.
fixed_priority_service_execution = true ∧
( ∀ b ∈ RTConfiguration. BandedConnection
∃ l ∈ RTConfiguration. Lanes :
    (b.low_range ≤ l.priority ≤ b.high_range)) ∧
RealTimeConfiguration.PriorityModelPolicy.
priority_model = "SERVER_DECLARED") ∨
(RTRequirement.
fixed_priority_service_execution = false ∧
RealTimeConfiguration.PriorityModelPolicy.
priority_model = "CLIENT_PROPOGATED")
```

If this expression evaluates to TRUE on an instance model, then it implies that the QoS configuration for this particular property has been mapped correctly. This applies to the `RTRequirements` and `RealTimeConfiguration` classes, and correspondence condition is added as a link between these classes in the metamodel. Similar to correspondence condition between `RTRequirements` and `RealTimeConfiguration` we described, other conditions for each of the QoS mapping rules have been specified ensuring that the transformation is verified correct if all these conditions are satisfied.

Note that the transformation verification does not imply the correctness of the QoS mapping rules themselves – rather it provides an assurance that the model transformation specification and implementation correctly mapped the QoS specifications that were formulated before. The validity of the generated QoS configuration is described in the next section.

### *4.2 Verifying the Generated QoS Configurations*

QUICKER uses a novel approach to check the correctness of the generated configurations. QUICKER achieves this capability by reducing the problem to a model checking problem. This section illustrates how the correctness of QoS configuration mappings is verified using the Bogor model-checking framework, which is a customizable explicit-state model checker implemented as an Eclipse plugin.

Verifying a system using Bogor involves defining (1) a model of the system using the *Bogor Input Representation* (BIR) language and (2) the *property* (*i.e.*, specification) that the model is expected to satisfy. Bogor then traverses the system model and checks whether or not the property holds. To validate QoS configuration options of an application using Bogor, we need to specify the application model and its QoS configurations. We use Bogor's extension features to customize the model-checker for resolving the QoS configuration challenges for component-based applications.

It is cumbersome to describe middleware QoS configuration options using the default input specification capabilities of BIR. This is because such a representation is at a much lower level of abstraction compared to domain-level concepts, such as components and QoS options, which we want to model-check. Additionally, specifying middleware QoS configuration options using BIR's low-level constructs can yield an unmanageably large state space since representing domain-level concepts with a low-level BIR specification requires additional auxiliary states that may be irrelevant to the properties being model-checked [23]. Therefore we have defined composite language constructs that represent functional sub-systems (such as components) and QoS options (such as thread pools) as though they were native BIR constructs.

Listing 1 shows an example of our QoS extensions in Bogor to represent QoS configuration options in middleware, which define two new data types: `Component`, which corresponds to a middleware component such as in our CIAO middleware, and `QoSOptions`, which captures QoS configuration options, such as `lane`, `band`, and `threadpool`.

In addition to defining constructs that represent domain concepts, such as components and QoS options, we also need to specify the *property* that the application should satisfy. In our case, a property simply means whether or not the QoS configurations are verified correct. Thus, since we need to verify the values of various QoS options as a means to check whether the application property is satisfied, we define *rules*

```
extension QoSOptions for
edu.ksu.cis.bogor.module.QoSOptions.QoSOptionsModule
{
  // Defines the new type to be used for
  typedef lane;
  typedef band;
  typedef threadpool;
  typedef prioritymodel;
  typedef policy;
  // Lane constructor.
  expdef QoSOptions.lane createLane (
   int static, int priority, int dynamic);
  // ThreadPool constructor.
  expdef QoSOptions.threadpool
  createThreadPool (boolean allowreqbuffering,
   int maxbufferedrequests, int stacksize, int
   maxbuffersize, boolean allowborrowing);
  // Set the band(s) for QoS policy.
  actiondef registerBands (QoSOptions.policy
   policy, QoSOptions.band ...);
  // Set the lane(s) for QoS policy.
  actiondef registerLanes (QoSOptions.policy
   policy, QoSOptions.lane ...);
  ...
}
extension Quicker for
edu.ksu.cis.bogor.module.Quicker
{
  // Defines the new type.
  typedef Component;
  // Component Constructor.
  expdef Quicker.Component
  createComponent (string component);
  // Set the QoS policy for the component.
  actiondef registerQoSOptions (Quicker.Component
    component,QoSOptions.policy policy);
  // Make connections between components.
  actiondef connectComponents (Quicker.Component
    server,Quicker.Component client);
  ...
}
```

Listing 1: **QUICKER BIR extension**

that capture values of these QoS options. BIR primitives are used to express these rules in the input specification of DRE system. Primitives are also used to capture component interconnections in BIR format which are required for populating the dependency structure for the specified input application. They are also used later during verification of options for connected components. We demonstrate this capability in the context of a representative system in Section 5.2.

Applications that need to be model-checked by Bogor must be represented in BIR format. Writing and maintaining BIR manually can be tedious and error-prone for domain experts (*e.g.*, avionics engineers) since configuring application QoS policies is typically done iteratively. Depending on the number of components and available configuration options, manual processes do not scale well.

To automate the process of creating BIR specification of applications, we therefore used the generative capabilities in GME to automatically generate BIR specification of an application from its QoS configurations model. This generative process is done in GME using a model interpreter that traverses the QoS configurations model and generates a BIR file that captures the application structure and its QoS properties. Our toolchain therefore automates the entire process of mapping application QoS policies to middleware QoS options, as well as converting these QoS options into BIR. A second model interpreter is used to generate the Real-time CCM-specific descriptors required to configure functional and QoS properties of an application and deploy it in its target environment. In the next section we empirically validate these generated QoS configurations.

### *4.3 Resolution of Challenges 2, 3 & 4*

Challenge 2 from Section 2.2 is resolved by the assurance of correctness provided by structural correspondence as follows. Pre- and post-conditions can be satisfied by assuring that the generated configurations are indeed the expected artifacts. Since the transformations are automated, developers need not worry about the plethora of configuration options and their interactions. Not all invariants may be satisfied by the transformation process alone. For example, whether the ORB can eliminate or minimize priority inversions depends solely on the quality of the middleware implementation and the OS platform on which it runs. To address this shortcoming, a potential solution may require that the transformation algorithms be enhanced to incorporate features of the available platform-specific support, such as the real-time features provided by the OS.

Finally, the model-checking extensions ensure that the system QoS configuration is valid at the local- and global-level including all the dependencies. This addresses Challenge 3. The entire process is automated in QUICKER and thus, can be repeated when changes in QoS policies of an application occur. For example, QUICKER helps DRE systems to evolve by automatically (re-)calculating the dependencies between options, and can thus be used repeatedly by DRE developers during the entire lifecycle thereby addressing Challenges 3 and 4.

Note that QoS configuration challenges also arise during runtime. We do not expect model checking to be applied at runtime. Rather that in the context of DRE systems, system developers will formally define the regions of operation of the system and the desird QoS policies as system conditions vary. Accordingly QUICKER can synthesize and verify the configurations for all the operating regions. The middleware is then subsequently responsible for adapting to the new configurations as the operating region of the system changes. QUICKER guarantees a priori availability of QoS configurations that are tailored to each operating region.

## 5 Validating QUICKER-generated QoS Configurations

This section validates the effectiveness of the QUICKER transformation algorithms in meeting the QoS requirements of DRE systems. We demonstrate these validations in the context of a DRE system as a proof-of-concept. Extensive user studies are necessary to provide more validation of the approach's effectiveness.

First we present a representative DRE system case study we used for the evaluation. We then describe how we have applied Bogor model-checking tool for verification of the generated configurations. Finally, through empirical evaluation, we validate the generated QoS options.

### 5.1 DRE system case study

The Basic Single Processor (BasicSP) is a scenario from the Boeing Bold Stroke component avionics computing product line. BasicSP uses a publish/subscribe service for event-based communication among its components, and has been developed using a QoS-enabled component middleware platform. The application is deployed using a single deployment plan on two physical nodes.

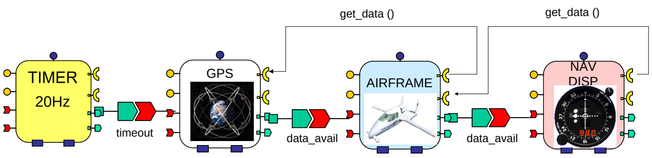


Fig. 11: Basic single processor

A GPS device sends out periodic position updates to a GUI display that presents these updates to a pilot. The desired data request and the display frequencies are at 20 Hz. The scenario shown in Figure 11 begins with the *GPS* component being invoked by the *Timer* component. On receiving a pulse event from the *Timer*, the *GPS* component generates its data and issues a data available event. The *Airframe* component retrieves the data from the *GPS* component, updates its state, and issues a data available event. Finally, the *NavDisplay* component retrieves the data from the *Airframe* and updates its state and displays it to the pilot. In its normal mode of operation, the *Timer* component generates pulse events at a fixed priority level, although its real-time configuration can be easily changed such that it can potentially support multiple priority levels.

It is necessary to carefully examine the end-to-end application critical path and configure the system components correctly such that the display refresh rate of 20 Hz may be satisfied. In particular, the latency between *Airframe* and *NavDisplay* components needs to be minimized to achieve the desired end goal. To this end, several characteristics of the BasicSP components are important and must be taken into account in determining the most appropriate QoS configuration space. For example, the *NavDisplay* component receives update events only from the *Airframe* component and does not send messages back to the sender *i.e.*, it just plays the role of a client. The *Airframe* component on the other hand communicates with both the *GPS* and *NavDisplay* components thereby playing the role of a client as well as a server. Various QoS options provided by the target middleware platform (in case of BasicSP, it is RT-CCM) ensure that these application level QoS requirements are satisfied. In the remainder of the paper, we focus on verification and validation of the QoS options generated using our approach.

### 5.2 Validating BasicSP QoS Configuration Dependencies via Model-checking

The QoS extensions to the BIR format described in Section 4.2 are used in maintaining and resolving dependencies between application components. For example, consider a real-time configuration of BasicSP scenario in which each of the *GPS*, *AirFrame*, and *NavDisplay* components are configured to have `priority bands` for separate service invocation priorities and the *Timer* component is configured to support multiple priority levels during generation of pulse events. Given such a configuration, we have that `priority band` values at *GPS* (client) component must match `ThreadPoolLanes` at *Timer* (server) component *i.e.*, a direct configuration dependency exists between these two components.

Further, since the pulse events are subsequently reported to *AirFrame* and *NavDisplay* components there is a similar indirect dependency between `band` values at these components and `lanes` at *Timer* component. The dependency structure of BasicSP scenario is maintained in QoS extensions to track such dependencies between QoS options.

Figure 12 represents the dependency structure generated using QoS extensions with the given configurations for our BasicSP scenario. Occurrences of change in configurations of either of the dependent components are followed by detection of potential mismatches such that all dependencies are exposed and resolved during application QoS design iterations.

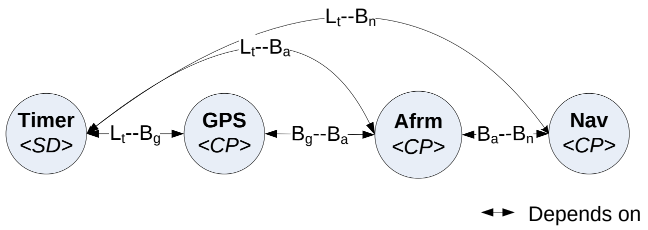


Fig. 12: **Dependency structure of BasicSP. $L_c$ denotes threadpool lane and $B_c$ denotes priority bands at component $c$. SD and CP indicate the SERVER_DECLARED and CLIENT_PROPAGATED priority models, respectively.**

A similar approach can be carried out if the DRE system has multiple different modes of operation, which in turn define multiple different regions of operation. Each such operating region may then have a set of configuration options. The Bogor-based model checking approach can now verify the dependencies among options for connected components for multiple modes of operation. The verified sets of configurations can then be used by the middleware to seamlessly adapt from one set to another depending on the mode.

### 5.3 Empirically Evaluating BasicSP QoS Configurations

In this section we empirically validate the effectiveness of the generated QoS configurations for the BasicSP case study.

**Experiment Configuration.** We have used ISISLab (`www.dre.vanderbilt.edu/ISISlab`) for evaluating observed QoS properties of DRE systems based on middleware QoS configurations generated using our configuration process. Each of the physical nodes used in our experiments was a 2.8 GHz Intel Xeon dual processor, 1 GB physical memory, 1 GHz network interface, and 40GB hard disks. Version 0.6 of our RT-CCM middleware CIAO was used running on Redhat Fedora Core release 4 with real-time preemption patches. The processes that hosted BasicSP components were run in the POSIX scheduling class SCHED_FIFO, enabling first-in-first-out scheduling semantics based on the priority of the process.

Table 1: **Generated QoS configuration for BasicSP**

| QoS configuration | Timer | GPS | Airframe |
|---|---|---|---|
| PriorityModel | SD | CP | CP |
| ThreadPool | | | |
| stacksize | 1024 | 1024 | 1024 |
| max_buff_reqs. | – | 20 | 20 |
| allow_borro. | FALSE | FALSE | FALSE |
| allow_req_buff | FALSE | TRUE | TRUE |
| Lane | | | |
| static_thrds | 4 | 8 | 8 |
| dyna_thrds | 0 | 0 | 0 |

As the first step, we modeled BasicSP QoS requirements using the requirements DSML described in Section 3. `bursty_client_requests` was set to FALSE for all components and `fixed_priority_service_execution` attribute was set to FALSE for every component except *Timer*. Secondly, we applied our model transformation algorithm to the requirements model above for generating detailed application configurations. Table 1 captures some of the important QoS configurations generated in our process. These configurations are represented as an application model. In the final step, we apply model interpreters for synthesizing descriptors required to configure the functional and QoS properties of the application during deployment.

In evaluating effectiveness of our configuration process, we collected end-to-end latency measurements between *Timer* and *NavDisplay* components. Earlier in Section 4.2 we discussed how correctness of QoS options can be verified using our process, the first experiment discussed below empirically evaluates the effectiveness of these options in meeting 20Hz operational display refresh rate of BasicSP from low to high workload conditions. Further, operational conditions of DRE system might change (unfavorably) during its execution. In order to evaluate the tolerance of our generated configurations under such conditions, in the second experiment we measure the metrics discussed above when invocation rate is steadily increased. Each of these experiments were performed for a constant time period and after executing 10,000 warmup iterations.

**Experiment 1: Increasing System Workload.** Figure 13 plots the latency measurements under increasing system workload. The workload is characterized [21] as a function performed with every client invocation. The signature of the function is given as: $void\ \ work(int\ units);$ where *units* argument specifies the amount of processor intensive work performed per call. The experiment was run for workload values of 10 through 80. End-to-end latency was observed to be at an average value of ∼1925 as can be seen in Figure 13a. Further, successive event-driven computations in the scenario exhibit an almost constant time complexity, indicated by relatively small dispersion in latencies as plotted in Figure 13b.

**Experiment 2: Increasing System Invocation Rate.** Performance of the generated configurations for BasicSP is given in Figure 14. Throughout this experiment the rate of invocation was increased from a normal operational value of 20Hz to a maximum of 40000Hz. Latency results are shown in Figure 14a which plots maximum, mean and minimum delay measurements for each invocation rate data point. Even with increasing rate the mean latency did not change significantly and was observed to be consistently just above 1900 microsec. for the entire range of invocation rates. Note that this is a desirable characteristic since even with an unfavorable change in operational conditions (*i.e.*, change in invocation rate) the latency was observed to be constant.

Jitter in latencies for each invocation rate is plotted in Figure 14b which shows that the deviation is bound between a high value of 42.44 (at 40Hz) and low value of 26.68 (at 2500Hz). At frequencies 2500Hz and higher the jitter values became quite stable showing a maximum variation of only 2.11. Overall, our results indicate that even under increased rate of invocation, the configurations perform effectively in achieving BasicSP latency requirements.

## 6 Related Work

QUICKER is designed to bridge the gap shown in Figure 15 between:

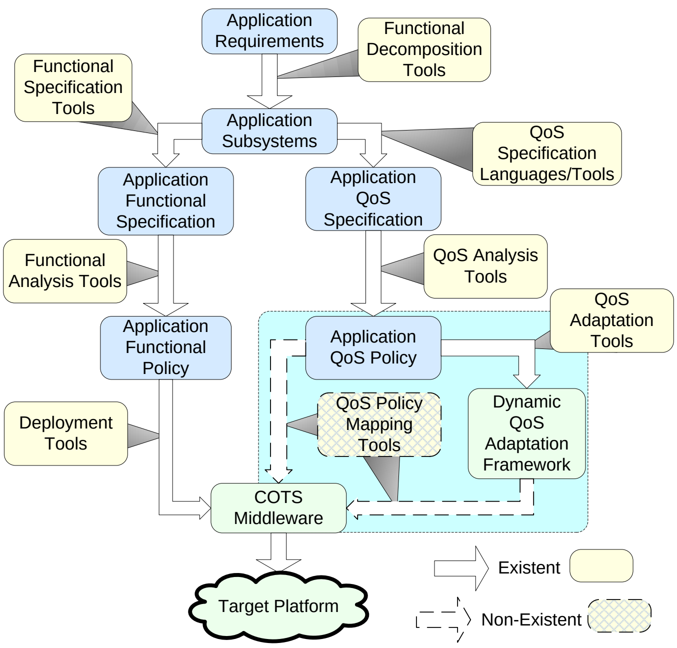


Fig. 15: QoS mapping landscape

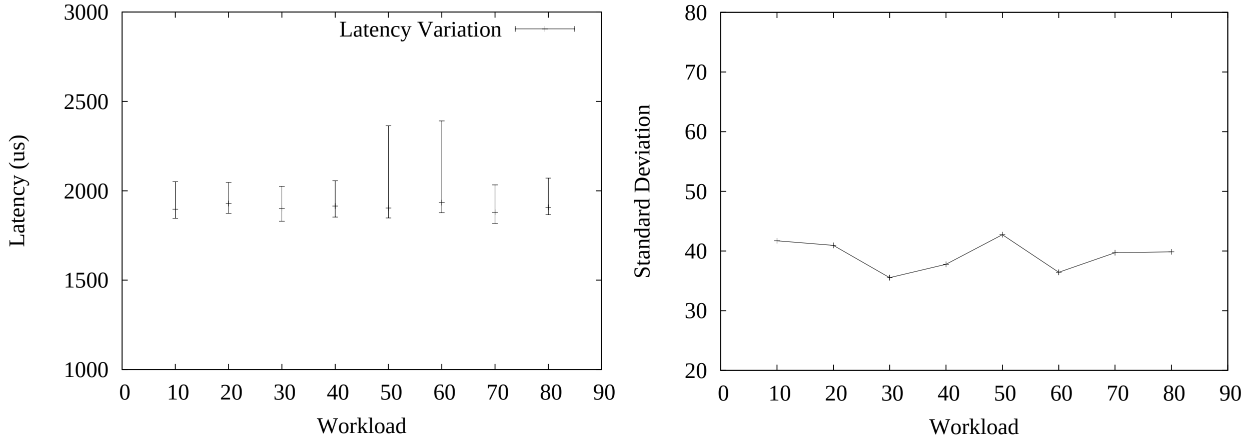


(a) End-to-end latency measurements (b) Statistical dispersion of latency

Fig. 13: **Evaluating BasicSP QoS configurations against increasing workload at a constant 20Hz invocation rate.**

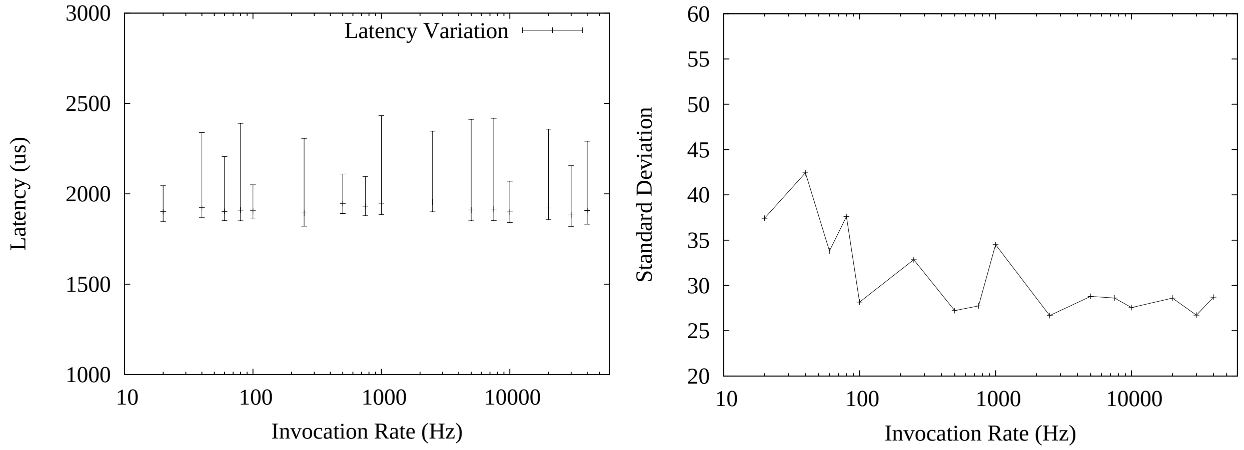


(a) End-to-end latency measurements (b) Statistical dispersion of latency

Fig. 14: **Evaluating BasicSP QoS configurations against increasing invocation rate: All the plots use logarithmic X axis and linear Y axis.**

- **Functional specification and analysis tools**, such as PICML and Cadena [7], that allow specification and analysis of application structure and behavior,
- **Schedulability analysis tools**, such as TIMES [2], AIRES [14], VEST [25], that perform schedulability and timing analysis to determine the exact priorities and time periods for application components, and
- **Dynamic QoS adaptation frameworks**, such as the Resource Adaptation and Control Engine (RACE) [24] and QuO [27], that allocate resources to application components, monitor the QoS of the system continuously, and apply corrective control to modify the QoS configuration of the middleware at runtime.

We present related work comparing and contrasting them with QUICKER.

**Validation and Analysis Techniques.** Model-driven techniques in [15,26] rely on a visual interface to help developers select a wide array of middleware QoS options for their applications. Such information is later used for generating test suites for purposes of empirical evaluation. In contrast, our configuration process does not expose the developers to all of the configuration space of underlying middleware and relies on platform-specific heuristics for generating QoS configurations. Further, using our process, the correctness of generated configurations is established in the design time. We argue that since our transformation algorithms codify best practices and patterns in middleware QoS configuration, QoS design and evolution throughout the system lifecycle using our approach is faster.

Analysis tools such as VEST [25], Cadena [7] and AIRES [14] evaluate whether certain timing, memory, power, and cost constraints and functional dependencies of real-time and embedded applications are satisfied. Our configuration process can be used as a complementary QoS design and analysis technique to these efforts since it emphasizes on mechanisms to (1) translate design-intent into actual configuration options of underlying middleware and (2) verify that both the transformation and subsequent modifications to the configuration options remain semantically valid.

**QoS Design and Specification Techniques.** The Adaptive Quality Modeling Language (AQML) [18] provides QoS adaptation policy modeling artifacts. AQML generators can (1) translate the QoS adaption policies (specified in AQML) into Matlab Simulink/Stateflow models for simulations using a control-centric view of QoS adaptation and (2) generate Contract Definition Language (CDL) specifications from AQML models to be used in target middleware. Our work differs with AQML since middleware model in QUICKER precisely abstracts the actual real-time CORBA implementation and does not need a two-level declarative translation (from AQML to CDL to target middleware) to achieve QoS configuration.

Ritter *et.al.* [22] describe CCM extensions for generic QoS support and discuss a QoS metamodel that supports domain-specific multi-category QoS contracts. The work in [8] focuses on capturing QoS properties in terms of *interaction patterns* among system components that are involved in executing a particular service and supporting run-time monitoring of QoS properties by distributing them over components (which can be monitored) to realize that service. Another approach that uses an aspect-oriented specification technique for component-based distributed systems is discussed in [4]. This work deals with specification of functional behavior, non-functional behavior, QoS management policies, and requirements of the application and synthesis of QoS management components for that supporting application-level adaptation strategies.

In contrast to the projects and tools described above, our work focuses on automating the error-prone activity of mapping platform-independent QoS policies to middleware-specific QoS configuration options. Representing QoS policies as model elements allows for a unified (with functional aspects of the application) and flexible QoS specification mechanism, while automating evolution of the QoS policies with application evolution; the platform-independent QoS policies also allow configurable re-targeting of the QoS mapping to support other types of middleware technologies.

## 7 Concluding Remarks

With the trend towards implementing key DRE system infrastructure at the middleware level, achieving the desired QoS is increasingly becoming more of a configuration problem than a development problem. The flexibility of configuration options in QoS-enabled component middleware, however, has created a new set of challenges. Key challenges include determining the correct set of values for the configuration options, understanding the dependency relations between the different options, and evolving the QoS configurations with changes to application functionality.

In this paper, we discussed how model transformations provided by our QUICKER tool automates the mapping of DRE system QoS policies into middleware-specific QoS configuration options. We showed how structural correspondence between input and output languages in our model-driven approach can be used to establish that initial system requirements are correctly mapped to middleware QoS options. We verified the correctness of generated QoS options using a model-checker and empirically showed that they are effective in satisfying system requirements.

We demonstrated the ideas behind QUICKER concretely in the context of the CIAO Lightweight CORBA Component Model middleware. By combining model transformation and generative techniques with advanced model-checking technologies, QUICKER automates the mapping of QoS policies of applications to QoS configuration options for the CIAO middleware technology. However, QUICKER's separation of platform-independent and platform-dependent concerns enables the use of PICML models, which is our source DSML, to specify QoS policies that can be mapped to other types of middleware, such as Web Services and EJB.

As a result, developers can concentrate on inherent complexities in the application domain rather than wrestle with low-level middleware-specific configuration options. QUICKER also helps ensure the validity of the values for the QoS configuration options, both at the individual component (local) level and at the aggregate application (global) level.

The following is a summary of lessons learned from our experience in using QUICKER:

• **QoS mapping is critical to successful deployment of systems built using component middleware.** With the increase in configuration complexity, the QoS mapping capabilities provided by QUICKER are essential to managing the complexity. Configuration of middleware options to achieve the desired QoS in DRE systems can be viewed as an directed acyclic graph whose root is the high-level mission requirements, edges are the individual mappings joining the vertices in a top-down fashion, and the vertices correspond to the different options available at each intermediate layer of abstraction. QUICKER is a part of a chain of mappings starting from high-level mission requirements to the actual deployment platform, and resides between the application components and the underlying component middleware implementation.

By employing DSMLs, QUICKER not only simplifies the QoS mapping process for DRE system developers, it also preserves the rich semantics associated with the mapping between the QoS policies and QoS configuration options at this level. Using such tools also helps QUICKER integrate with mapping technologies that exist both *above* (*e.g.*, mission requirement mapping tools, functional decomposition tools, and functional analysis tools) and *below* (*e.g.*, deployment planning tools) the level at which QUICKER operates in a component-based DRE system development lifecycle.

• **Horizontal mapping of QoS is as important as vertical QoS mapping.** QUICKER currently focuses on mapping application QoS policies onto a single underlying middleware technology. Large-scale DRE systems—particularly systems requiring dynamic resource management [16]—are often composed of heterogeneous technologies. It is therefore essential for QoS mapping tools to not only support *vertical mapping* (*i.e.*, the mapping of policies and validation onto a single technology) but also *horizontal mapping* (*i.e.*, the mapping of QoS policies onto multiple heterogeneous technologies, while reconciling the differences between these technologies). Until such mapping is performed, QoS configuration and associated tools will remain as islands, which significantly complicates integration efforts for large-scale DRE systems. Care must be taken, however, to not introduce additional learning curve for developers.

In the future in order to show its scalability we plan to apply and evaluate our technique on complex and large-scale DRE systems. Our current approach is one-dimensional *i.e.*, both the QoS requirements mapping and configuration validation is done for a single dimension (such as real-time request-response or publish-subscribe communication dimensions). In the future we plan to investigate and develop configuration techniques under simultaneous requirements across distinct QoS dimensions. An effort is underway in extending our process for other component middleware platforms that exhibit the same level of configurability. As part of this effort, we are looking at development of parameterized model transformations that allow specification of templatized QoS mappings and later generation of platform-specific QoS mapping instances by specializing these templatized mappings. These parametrized models are, however, unlike the intermediate representations used in related efforts, such as AQML [18].

The QUICKER tool is available as open-source from `www.dre.vanderbilt.edu/CoSMIC/`.

## References


1. Ákos Lédeczi, Árpád Bakay, M. Maróti, P. Völgyesi, G. Nordstrom, J. Sprinkle, and G. Karsai. Composing Domain-Specific Design Environments. *Computer*, 34(11):44–51, 2001.
2. T. Amnell, E. Fersman, L. Mokrushin, P. Pettersson, and W. Yi. TIMES: A Tool for Schedulability Analysis and Code Generation of Real-Time Systems. In K. G. Larsen and P. Niebert, editors, *Formal Modeling and Analysis of Timed Systems: First International Workshop, FORMATS 2003, Marseille, France, September 6-7, 2003. Revised Papers*, volume 2791 of *Lecture Notes in Computer Science*, pages 60–72. Springer, 2003.
3. K. Balasubramanian, J. Balasubramanian, J. Parsons, A. Gokhale, and D. C. Schmidt. A Platform-Independent Component Modeling Language for Distributed Real-Time and Embedded Systems. In *RTAS '05: Proceedings of the 11th IEEE Real Time on Embedded Technology and Applications Symposium*, pages 190–199, Washington, DC, USA, 2005. IEEE Computer Society.
4. L. Blair, G. S. Blair, A. Anderson, and T. Jones. Formal Support For Dynamic QoS Management in the Development of Open Component-based Distributed Systems. *IEEE Software*, 148(3), Nov. 2001.
5. K. Czarnecki and S. Helsen. Feature-based Survey of Model Transformation Approaches. *IBM Syst. J.*, 45(3):621–645, 2006.
6. G. Deng, C. Gill, D. C. Schmidt, and N. Wang. QoS-enabled Component Middleware for Distributed Real-Time and Embedded Systems. In I. Lee, J. Leung, and S. Son, editors, *Handbook of Real-Time and Embedded Systems*. CRC Press, 2007.
7. J. Hatcliff, W. Deng, M. Dwyer, G. Jung, and V. Prasad. Cadena: An Integrated Development, Analysis, and Verification Environment for Component-based Systems. In *Proceedings of the 25th International Conference on Software Engineering*, pages 160–172, Portland, OR, May 2003.
8. Jaswinder Ahluwalia and Ingolf H. Krger and Walter Phillips and Michael Meisinger. Model-Based Run-Time Monitoring of End-to-End Deadlines. In *Proceedings of the Fifth ACM International Conference On Embedded Software*, pages 100–109, Jersey City, NJ, Sept. 2005. ACM.
9. G. Karsai, A. Agrawal, F. Shi, and J. Sprinkle. On the Use of Graph Transformation in the Formal Specification of Model Interpreters. *Journal of Universal Computer Science*, 9(11):1296–1321, 2003. `www.jucs.org/jucs_9_11/on_the_use_of`.
10. A. Kavimandan, K. Balasubramanian, N. Shankaran, A. Gokhale, and D. C. Schmidt. Quicker: A model-driven qos mapping tool for qos-enabled component middleware. In *ISORC '07: Proceedings of the 10th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing*, pages 62–70, Washington, DC, USA, 2007. IEEE Computer Society.
11. A. Kavimandan and A. Gokhale. Automated Middleware QoS Configuration Techniques using Model Transformations. In *Proceedings of the* $14^{th}$ *IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2008)*, pages 93–102, St. Louis, MO, USA, Apr. 2008.
12. A. Kavimandan, A. Narayanan, A. Gokhale, and G. Karsai. Evaluating the Correctness and Effectiveness of a Middleware QoS Configuration Process in Distributed Real-time and Embedded Systems. In *Proceedings of the* $11^{th}$ *IEEE International Symposium on Object-oriented Real-time distributed Computing (ISORC 2008)*, pages 100–107, Orlando, FL, USA, May 2008.
13. A. Kleppe, J. Warmer, and W. Bast. *MDA Explained: The Model Driven Architecture*($MDA^{TM}$)*: Practice and Promise*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, Apr 2003.
14. S. Kodase, S. Wang, Z. Gu, and K. G. Shin. Improving Scalability of Task Allocation and Scheduling in Large Distributed Real-time Systems using Shared Buffers. In *Proceedings of the 9th Real-time/Embedded Technology and Applications Symposium (RTAS 2003)*, Washington, DC, May 2003. IEEE.
15. A. S. Krishna, E. Turkay, A. Gokhale, and D. C. Schmidt. Model-Driven Techniques for Evaluating the QoS of Middleware Configurations for DRE Systems. In *Proceedings of the 11th Real-time Technology and Application Symposium (RTAS '05)*, pages 180–189, San Francisco, CA, Mar. 2005. IEEE.
16. P. Lardieri, J. Balasubramanian, D. C. Schmidt, G. Thaker, A. Gokhale, and T. Damiano. A Multi-layered Resource Management Framework for Dynamic Resource Management in Enterprise DRE Systems. *Journal of Systems and Software: Special Issue on Dynamic Resource Management in Distributed Real-time Systems*, 80(7):984–996, July 2007.
17. A. Narayanan and G. Karsai. Verifying Model Transformations by Structural Correspondence. Technical Report ISIS-07-809, Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN, Dec 2007.

18. S. Neema, T. Bapty, J. Gray, and A. Gokhale. Generators for Synthesis of QoS Adaptation in Distributed Real-time Embedded Systems. In *Proceedings of the ACM SIGPLAN/SIGSOFT Conference on Generative Programming and Component Engineering (GPCE'02)*, pages 236–251, Pittsburgh, PA, Oct. 2002.
19. Object Management Group. *Lightweight CCM RFP*, realtime/02-11-27 edition, Nov. 2002.
20. Object Management Group. *Real-time CORBA Specification*, 1.2 edition, Jan. 2005.
21. I. Pyarali, D. Schmidt, and R. Cytron. Techniques for enhancing real-time corba quality of service. *Proceedings of the IEEE*, 91(7):1070–1085, July 2003.
22. T. Ritter, M. Born, T. Unterschütz, and T. Weis. A QoS Metamodel and its Realization in a CORBA Component Infrastructure. In *Proceedings of the* $36^{th}$ *Hawaii International Conference on System Sciences (HICSS'03)*, page 318, Honolulu, HI, Jan. 2003.
23. Robby, M. Dwyer, and J. Hatcliff. Bogor: An Extensible and Highly-Modular Model Checking Framework. In *Proceedings of the* $4^{th}$ *Joint Meeting of the European Software Engineering Conference and ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE 2003)*, Helsinki, Finland, September 2003. ACM.
24. N. Shankaran, D. C. Schmidt, Y. Chen, X. Koutsoukous, and C. Lu. The Design and Performance of Configurable Component Middleware for End-to-End Adaptation of Distributed Real-time Embedded Systems. In *Proc. of the 10th IEEE International Symposium on Object/Component/Service-oriented Real-time Distributed Computing (ISORC 2007)*, Santorini Island, Greece, May 2007.
25. J. A. Stankovic, R. Zhu, R. Poornalingam, C. Lu, Z. Yu, M. Humphrey, and B. Ellis. VEST: An Aspect-Based Composition Tool for Real-Time Systems. In *RTAS '03: Proceedings of the The 9th IEEE Real-Time and Embedded Technology and Applications Symposium*, page 58, Washington, DC, USA, 2003. IEEE Computer Society.
26. L. Zhu, N. B. Bui, Y. Liu, and I. Gorton. MDABench: Customized benchmark generation using MDA. *Journal of Systems and Software*, 80(2):265–282, Feb. 2007.
27. J. A. Zinky, D. E. Bakken, and R. Schantz. Architectural Support for Quality of Service for CORBA Objects. *Theory and Practice of Object Systems*, 3(1):1–20, 1997.