

# Design and Performance Evaluation of Configurable Component Middleware for End-to-End Adaptation of Distributed Real-time Embedded Systems

Nishanth Shankaran<sup>†</sup>, Douglas C. Schmidt<sup>†</sup>, Xenofon D. Koutsoukos<sup>†</sup>,

Yingming Chen<sup>‡</sup>, and Chenyang Lu<sup>‡</sup>

<sup>†</sup>Dept. of EECS

<sup>‡</sup>Dept. of Computer Science and Engineering,

Vanderbilt University, Nashville, TN

Washington University, St. Louis

**Abstract**—Standards-based quality of service (QoS)-enabled component middleware is increasingly being used as a platform for developing distributed real-time embedded (DRE) systems that execute in open environments where operational conditions, input workload, and resource availability cannot be characterized accurately a priori. Although QoS-enabled component middleware offers many desirable features, until recently it lacked the ability to efficiently allocate resources and configure platform-specific QoS settings based on utilization of system resources and application QoS. Moreover, it has also lacked the ability to monitor and enforce application QoS requirements.

This paper presents two contributions to research on adaptive resource management for component-based DRE systems. First, we describe the structure and functionality of the Resource Allocation and Control Engine (RACE), which is an open-source adaptive resource management framework built atop standards-based QoS-enabled component middleware. Second, we demonstrate and evaluate the effectiveness of RACE in the context of a representative DRE system: NASA’s Magnetospheric Multi-scale Mission system. Our empirical results show that the capabilities provided by RACE yields a predictable and high performance system, even in the face of changing operational conditions, workloads, and resource availability.

## I. INTRODUCTION

**Emerging trends and challenges.** *Distributed real-time and embedded* (DRE) systems form the core of many mission-critical domains, such as shipboard computing environments, avionics mission computing, multi-satellite missions, and intelligence, surveillance and reconnaissance missions. *Quality of service (QoS)-enabled distributed object computing (DOC) middleware* based on standards like Real-time CORBA (RT-CORBA) [1] and the Real-Time Specification for Java (RTSJ) [2] have been used to develop such DRE systems. More recently, *QoS-enabled component middleware*, such as the Lightweight CORBA Component Model (CCM) [3] and PRiSm [4], have been used as the middleware for DRE systems [5].

Compared to QoS-enabled DOC middleware, QoS-enabled component middleware capabilities enhance the design, development, evolution, and maintenance of DRE systems. Examples of additional capabilities offered by QoS-enabled component middleware include (1) standardized interfaces for application component interaction, (2) model-based tools for deploying and interconnecting components, and (3) standards-based mechanisms for installing, initializing, and configuring application components, thus separating concerns of application development, configuration, and deployment.

In prior work, we developed a domain-specific modeling language (DSML) called the *Platform-Independent Component Modeling Language* (PICML) [6] that alleviates many complexities associated with developing component-based

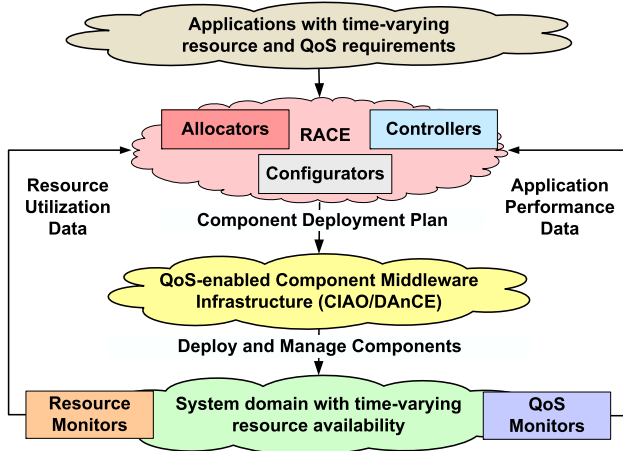
DRE systems using the *Component-Integrated ACE ORB* (CIAO) [7]. CIAO abstracts key Real-time CORBA QoS concerns (such as priority models, thread-to-connection bindings, and timing properties) into elements that can be configured declaratively via Lightweight CCM metadata (such as standards for specifying, implementing, packaging, assembling, and deploying components). PICML enables DRE system developers to (1) design component interfaces and compose applications by interconnecting components, (2) specify QoS and resource utilization characteristics of applications such as end-to-end deadlines, *estimated* CPU, memory, and network bandwidth utilization characteristics, (3) configure middleware, operating system, and network QoS parameters, (4) specify *estimated* resource availability of the DRE system, (5) allocate resource to components that make up the application, and (6) generate deployment metadata used by component middleware to deploy applications.

Although CIAO and PICML raise the level of abstraction used to develop DRE systems relative to DOC middleware, unresolved challenges remain. In particular, when DRE system developers allocate resources to, and configure QoS settings of, applications using PICML, these operations are performed based on *estimated* resource utilization of applications and *estimated* availability of system resources. These estimates may be imprecise for open DRE systems that execute in environments where operational conditions, input workload, and resource availability cannot be characterized accurately a priori.

In general, there are limited mechanisms in existing QoS-enabled component middleware platforms and DSML tools to (1) specify end-to-end QoS requirements and (2) monitor application behavior to ensure that these QoS requirements are met. Moreover, when applications are composed at runtime by intelligent mission planners [8], only end-to-end QoS requirements of the applications are specified. What is needed, therefore, are middleware-centric capabilities for allocating resources automatically and monitoring/(re)configuring QoS settings of applications to enforce their end-to-end QoS requirements.

**Solution: A component-based adaptive resource management framework.** To address the needs of DRE system developers outlined above, we have developed the *Resource Allocation and Control Engine* (RACE), which is an adaptive resource management framework built atop our CIAO QoS-enabled component middleware. As shown in Figure 1, RACE provides (1) *resource monitor* components that track utilization of various system resources, such as CPU, memory, and network bandwidth, (2) *QoS monitor* components that

track application QoS, such as end-to-end delay, (3) *resource allocator* components that allocate resource to components based on their resource requirements and current availability of system resources, (4) *configurator* components that configure QoS parameters of application components, (5) *controller* components that compute end-to-end adaptation decisions to ensure that QoS requirements of applications are met, and (6) *effector* components that perform controller-recommended adaptations.



1: Resource Allocation and Control Engine (RACE) for DRE Systems

RACE supports multiple applications running in various DRE system environments and allows applications with diverse QoS requirements to share resources simultaneously. RACE’s allocator and controller entities can be configured with diverse resource allocation and control algorithms. It also uses standard component middleware deployment and configuration mechanisms [9] to allocate resources to applications and control DRE system performance once applications are deployed and running.

This paper provides two contributions to research on adaptive resource management for component-based DRE systems. First, it describes the component-based design of the RACE framework. Second, we qualitatively and quantitatively evaluate the effectiveness of RACE in resolving key adaptive resource management challenges of a representative DRE system.

The remainder of the paper is organized as follows: Section II compares our research on RACE with related work; Section III motivates the use of RACE in the context of a representative DRE system; Section IV describes the architecture of RACE and shows how it meets the QoS requirements of the DRE system described in Section III; Section V empirically evaluates how RACE improves the QoS of the DRE system and presents an empirical measure of the overhead associated with the RACE framework; and Section VI presents concluding remarks.

## II. RELATED WORK

This section compares our work on RACE with related research on building large-scale DRE systems. As shown below, we classify this research along two orthogonal dimensions: (1)

QoS-enabled DOC middleware vs. QoS-enabled component middleware and (2) design-time vs. run-time QoS configuration, optimization, analysis, and evaluation of constraints, such as timing, memory, and CPU.

### A. QoS-enabled DOC Middleware

**Design-time.** RapidSched [10] enhances QoS-enabled DOC middleware, such as RT-CORBA, by computing and enforcing distributed priorities. RapidSched uses PERTS [11] to specify real-time information, such as deadline, estimated execution times, and resource requirements. Static schedulability analysis (such as rate-monotonic analysis) is then performed and priorities are computed for each CORBA object in the system. After the priorities are computed, RapidSched uses RT-CORBA features to enforce these computed priorities.

**Run-time.** Early work on resource management middleware for shipboard DRE systems presented in [12], [13] motivated the need for adaptive resource management middleware. This work was further extended by QARMA [14], which provides resource management as a *service* for existing QoS-enabled DOC middleware, such as RT-CORBA. Kokyu [15] also enhances RT-CORBA QoS-enabled DOC middleware by providing a portable middleware scheduling framework that offers flexible scheduling and dispatching services. Kokyu performs feasibility analysis based on estimated worst case execution times of applications to determine if a set of applications is *schedulable*. Resource requirements of applications, such as memory and network bandwidth, are not captured and taken into consideration by Kokyu. Moreover, Kokyu lacks the capability to track utilization of various system resources as well as QoS of applications. To address these limitations, research presented in [16] enhances QoS-enabled DOC middleware by combining Kokyu and QARMA.

Our work on RACE extends this earlier work on QoS-enabled DOC middleware by providing an adaptive resource management framework for DRE systems built atop QoS-enabled component middleware. DRE systems built using RACE benefit from the additional capabilities offered by QoS-enabled component middleware compared to QoS-enabled DOC middleware, as described in Section I. Moreover, the elements of RACE are designed as CCM components, so RACE itself can be configured using DSML tools, such as PICML [6].

### B. QoS-enabled Component Middleware

**Design-time.** Cadena [17] is an integrated environment for developing and verifying component-based DRE systems by applying static analysis, model-checking, and lightweight formal methods. Like PICML, Cadena also provides a component assembly framework for visualizing and developing components and their connections. VEST [18] is a DSML that enables embedded system composition from component libraries and checks whether timing, memory, power, and cost constraints of real-time and embedded applications are satisfied. AIRES [19] is another DSML that provides the means to map design-time models of component composition with real-time requirements to run-time models that weave together timing and scheduling attributes.

These tools are similar to PICML and use *estimates*, such as estimated worst case execution time, estimated CPU, memory, and/or network bandwidth requirements. These tools are targeted for systems that execute in *closed* environments, where operational conditions, input workload, and resource availability can be characterized accurately *a priori*. Since RACE tracks and manages utilization of various system resources, as well as application QoS, it can be used in conjunction with these tools to build DRE systems that execute in open environments.

**Run-time.** QoS provisioning frameworks, such as QuO and Qoskets [20] help ensure desired performance of DRE systems built atop QoS-enabled DOC middleware and QoS-enabled component middleware, respectively. When applications are designed using Qoskets (1) resources are dynamically (re)allocated to applications in response to changing operational conditions and/or input workload and (2) application parameters are fine-tuned to ensure that allocated resource are used effectively. With this approach, however, applications are augmented explicitly at design-time with Qosket components, such as monitors, controllers, and effectors. This approach thus requires redesign and reassembly of existing applications built without Qoskets. When applications are generated at run-time (*e.g.*, by intelligent mission planners [8]), this approach would require planners to augment the applications with Qosket components, which may be infeasible since planners are designed and built to solve mission goals and to work atop any component middleware, not just CCM.

Compared with related work, RACE provides adaptive resource and QoS management capabilities in a more transparent and non-intrusive way. In particular, it allocates CPU, memory, and networking resources to application components and tracks and manages utilization of various system resources, as well as application QoS. In contrast to our own earlier work on QoS-enabled DOC middleware, such as FC-ORB [21] and HiDRA [22], RACE is a QoS-enabled component middleware framework that enables the deployment and configuration of feedback control loops in DRE systems.

In summary, RACE's novelty stems from its combination of design-time DSML tools and QoS-enabled component middleware run-time platforms. RACE can be used to deploy and manage component-based applications that are composed at design-time via the PICML [6] DSML, as well as at run-time the SA-POP [8] intelligent mission planner (described in Section III-A). Moreover, RACE's reusable entities, such as resource monitors, QoS monitors, and effectors, can be configured to incorporate a range of existing control algorithms, such as EUCON [23] and HySUCON [24], as well as future algorithms.

### III. MOTIVATING APPLICATION SCENARIO

We use the NASA's upcoming Magnetospheric Multi-scale (MMS) mission ([stp.gsfc.nasa.gov/missions/mms/mms.htm](http://stp.gsfc.nasa.gov/missions/mms/mms.htm)) as a motivating DRE system example to evaluate the effectiveness and performance of RACE. First, we present an overview of the MMS mission, followed by the resource and QoS management challenges involved in developing the MMS mission using QoS-enabled component middleware.

#### A. MMS Mission Overview

The goal of the MMS mission is to study the micro-physics of three fundamental plasma processes occurring in the earth's magnetosphere: magnetic reconnection, particle acceleration, and turbulence. MMS mission consists of a constellation of identical spacecrafts that maintain a specific formation while orbiting over region of scientific interest (ROI). Since the plasma processes are inherently transient (especially magnetic reconnection), MMS missions requires reactive on-board autonomy to enable the spacecraft to transition between three modes of operation: slow survey, fast survey, and burst. Slow survey mode is entered outside the ROI and enables minimal data acquisition (primarily for health monitoring). The fast survey mode is entered when the spacecrafts are within a ROI, which enables data acquisition for all payload sensors at a moderate rate. While in fast survey mode, the data from a subset of the payload sensors is analyzed on board to detect the likelihood of a transient plasma event. If any plasma activity is detected, all the spacecrafts enter the burst mode and all payload instruments acquire data at high rates.

To address the challenges associated with efficient operation of the different configurations/modes outlined above and the transitions between them, on board intelligent mission planners such as *spreading activation partial order planner* (SA-POP) [8] have been developed. SA-POP decomposes overall mission goal(s) into sets of applications that can be executed concurrently. SA-POP employs decision-theoretic methods and other AI schemes (such as hierarchical task decomposition) to decompose mission goals into navigation, control, data gathering, and data processing applications.

In addition to initial generation of applications, SA-POP incrementally generates new applications in response to changing mission goals and/or degraded performance reported by the mission and system monitors. These applications are classified into two classes (*important* and *best-effort*) based on the operation performed by the application. For example, applications that are responsible for the *guidance–navigation–control* of the spacecrafts belong to the *important* class, where as data analysis applications belong to the *best-effort* class.

#### B. Challenges of Developing the MMS Mission using QoS-enabled Component Middleware

As discussed in Section I, the use of QoS-enabled component middleware to develop DRE systems, such as the NASA MMS mission, significantly improves the design, development, evolution, and maintenance of these large-scale systems. In the absence of an adaptive resource management framework like RACE, however, several key challenges remain unresolved when using component middleware. Below we present the key resource and QoS management challenges associated with the MMS mission DRE system.

**Challenge 1: Efficient resource allocation to applications.** Applications generated by SA-POP are *resource sensitive*, *i.e.*, end-to-end QoS is reduced significantly if the required type and quantity of resources are not provided to the applications at the right time. System resources should therefore be allocated in a timely fashion to components of applications such that their resource requirements are met. In open DRE systems

like MMS, however, input workload affects utilization of system resources by, and QoS of, applications. These parameters of the applications may therefore vary significantly from their estimated values. Moreover, system resource availability, such as available network bandwidth, may also be time variant. A resource management framework like RACE should therefore support multiple resource allocation strategies to handle the needs of heterogeneous applications, which include guidance, navigation, control, data acquisition, data handling, and data analysis applications.

**Challenge 2: Configuring platform-specific QoS parameters.** The QoS of applications depend on various platform-specific real-time QoS configurations including (1) QoS configuration of the QoS-enabled component middleware such as priority model, threading model, and request processing policy, (2) operating system QoS configuration such as real-time priorities of the process(es) and thread(s) that host and execute within the components respectively, and (3) networks QoS configurations, such as `difserv` code-points of the component interconnections. Since these configurations are platform-specific, it is tedious and error-prone for system developers or SA-POP to specify them in isolation. An adaptive resource management framework like RACE should therefore provide abstractions that shield developers and/or SA-POP from low-level platform-specific details and define higher-level QoS specification models.

**Challenge 3: Monitoring end-to-end QoS and ensuring QoS requirements are met.** To meet the end-to-end QoS requirements of applications, an adaptive resource management framework like RACE must provide monitors that track QoS of applications at run-time. Although some QoS properties (such as accuracy, precision, and fidelity of the produced output) are application-specific, certain QoS (such as *end-to-end delay*) can be tracked by the framework transparently to the application. The framework should also provide hooks into which application specific QoS monitors can be configured. The framework should enable the system to *adapt* to dynamic changes, such as variations in operational conditions, input workload, and/or resource availability, and thereby ensure that QoS requirements of applications are not violated.

#### IV. STRUCTURE AND FUNCTIONALITY OF RACE

RACE is built atop of the QoS-enabled component middleware CIAO and DANCE, which are open-source implementations of the OMG Lightweight CCM [3], Deployment and Configuration (D&C) [25], and RT-CORBA [1] specifications, which are outlined in Sidebar 1.

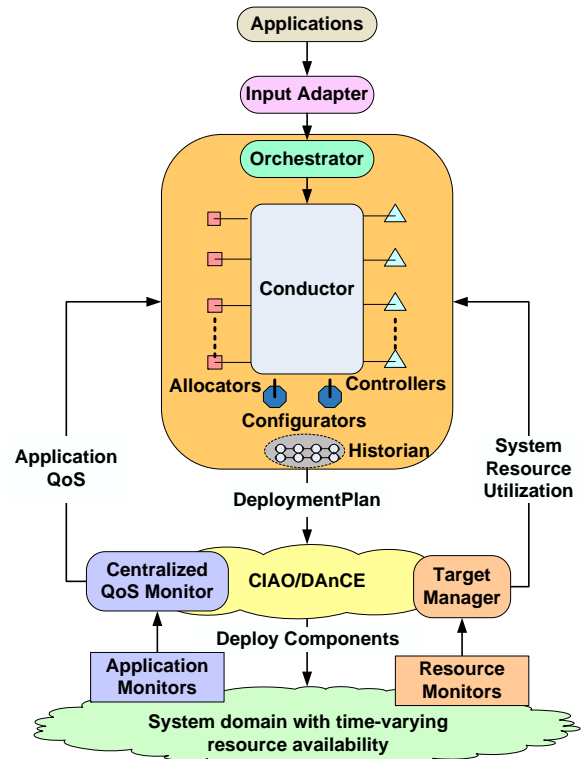
As shown in Figure 2, RACE is composed of the following components: (1) InputAdapter, (2) Orchestrator, (3) Conductor, (4) Allocators, (5) Controllers, (6) Configurators, and (7) Historian. RACE also monitors application QoS and system resource usage via its CentralizedQoSMonitor and TargetManager components. All components of RACE are deployed and configured using DANCE. This section motivates and describes the design of RACE by showing how it resolves the three challenges presented in the MMS case study from Section III.

### Sidebar 1: Overview of CCM Middleware

RT-CORBA adds QoS control to regular CORBA to improve application *predictability*, such as bounding priority inversions. RT-CORBA provides policies and mechanisms for configuring middleware features such as thread pools, priority models, protocol policies, and explicit binding. These capabilities address some, but by no means all, important DRE system development challenges. To address this issue, the OMG introduced the Lightweight CCM specification, which is built atop the RT-CORBA specification and standardizes the development, configuration, and deployment of component-based applications. Key entities of CCM-based component middleware include:

- **Component**, which encapsulates the behavior of the application. Components interact with clients and each other via *ports*, which are of four types (1) *facets*, (2) *receptacles*, (3) *event sources*, and (4) *event sinks*.
- **Container**, which provides an execution environment for components with common operating requirements. The container also provides an abstraction of the underlying middleware and enable the component to communicate via the underlying middleware bus and reuse common services offered by the underlying middleware.
- **Component server**, which is a generic server process that hosts application containers. One or more components can be collocated in one component server.

Since the Lightweight CCM specification does not standardize the process of deployment, initialization, and interconnection of components, the OMG Deployment and Configuration (D&C) specification was introduced as an addendum to the CCM specification. Our open-source implementation of the D&C specification *Deployment and Configuration Engine* (DAnCE) enables the deployment and configuration of components in DRE systems.



2: Structure and Interactions in RACE

### A. Efficient Resource Allocation

To allocate resources efficiently in an open DRE system, such as the NASA’s MMS mission system, RACE performs the following steps: (1) it parses the metadata that describes the application to obtain the resource requirement(s) of components that make up the application, (2) obtains current resource utilization from resource utilization monitors, and (3) selects and invokes an appropriate implementation(s) of resource allocation algorithm depending on the properties of the application and the overhead associated with the implementation(s). Below we describe the RACE components that work together to perform the steps outlined above and resolve the resource allocation challenges of the MMS mission as described in Section III-B.

**1. InputAdapter.** End-to-end applications can be composed in many ways. For example, an application can be composed by using a DSML like PICML at system design-time and/or by an intelligent mission planner like SA-POP at run-time. When an application is composed using PICML, metadata describing the application is captured in a XML file based on the `PackageConfiguration` schema defined by the D&C specification [25]. When applications are generated during runtime by SA-POP, metadata is captured in an in-memory structure defined by the planner.

RACE can be configured using a DSML such as PICML with appropriate `InputAdapter` that parses the metadata that describes the application into an in-memory end-to-end (E-2-E) IDL structure that is managed internally by RACE. The E-2-E IDL structure populated by the `InputAdapter` contains information regarding the application, including (1) components that make up the application and their resource requirement(s), (2) interconnections between the components, (3) application QoS properties (such relative priority) and QoS requirement(s) (such as end-to-end delay), and (4) mapping components onto domain nodes.<sup>1</sup>

**2. TargetManager.** As shown in Figure 2, RACE employs the `TargetManager` to obtain information regarding system resource utilization. `TargetManager` uses a hierarchical design and receives periodic resource utilization updates from `ResourceMonitors` within the domain. It uses these updates to track resource usage of all resources within the domain.

**3. Allocators** are implementations of resource allocation algorithms that allocate various domain resources (such as CPU, memory, and network bandwidth) to components of an application by determining the mapping of components onto nodes in the system domain. For certain applications—usually the important ones—*static* mapping between components and nodes may be specified at design-time by system developers. To honor these static mappings, RACE therefore provides a *static allocator* that ensures components are allocated to nodes in accordance with the static mapping specified in the application’s metadata. If no static mapping is specified, however, *dynamic allocators* determine the component to node

<sup>1</sup>The mapping of components onto nodes need not be specified in the metadata that describes the application which is given to RACE. If an mapping is specified, it is honored by RACE; if not, a mapping is determined at run-time by RACE’s `Allocators`.

mapping at run-time based on resource requirements of the components and current resource availability on the various nodes in the domain. Input to `Allocators` include the E-2-E IDL structure corresponding to the application and the current utilization of system resources. Since `Allocators` themselves are CCM components, RACE can be configured with new `Allocators` by using PICML.

The current version of RACE supports following algorithms as `Allocators`: (1) CPU allocator, (2) memory allocator, (3) network-bandwidth allocator, (4) PBFDD allocator [26] that allocates CPU, memory, and network-bandwidth, and (5) static allocator. Metadata is associated with each allocator and captures its type (*i.e.*, static, single dimension bin-packing [27], or PBFDD) and associated resource overhead (such as CPU and memory utilization).

**4. Orchestrator and Conductor.** After the metadata describing the application is parsed by RACE’s `InputAdapter`, the in-memory E-2-E IDL structure is passed onto the `Orchestrator`. This component processes the E-2-E structure to determine the types of resources (*e.g.*, CPU, memory, or network bandwidth) required and whether a static allocation is specified. If a static allocation is specified, the static allocator is selected; otherwise a dynamic allocator(s) is selected based on the type(s) of resources required. This selection process is captured in the `Composition` structure.

The `Orchestrator` passes the `Composition` and the E-2-E to the `Conductor`, which then performs the desired orchestration by invoking the `Allocator(s)` specified in the `Composition`, along with the resource utilization information obtained from the `TargetManager` to map components onto nodes in the system domain. After resources are allocated to the application, the `Conductor` converts the application from RACE’s internal E-2-E IDL structure into the standard `DeploymentPlan` IDL structure defined by the D&C specification [25]. The `DeploymentPlan` IDL structure is then passed to the underlying `DAnCE` middleware to deploy the components on the designated target nodes.

Since the elements of RACE are developed as CCM components, RACE itself can be configured using DSML tools, such as PICML. Moreover, new `InputAdapters` and `Allocators` can be plugged directly into RACE without modifying RACE’s existing architecture. RACE can be used to deploy and allocate resources to applications that are composed at design-time and run-time. RACE’s `Allocators` with inputs from the `TargetManager` allocates resource to application components based on runtime resource availability, thereby addressing the resource allocation challenge for DRE systems identified in Section III-B.

### B. QoS Parameter Configuration

RACE shields application developers and SA-POP from low-level platform-specific details and defines a higher-level QoS specification model. Developers and/or SA-POP specify only QoS characteristics of the application, such as QoS requirements and relative importance, and RACE automatically configures platform-specific parameters accordingly. Below, we describe the RACE components that work together to provide these capabilities and resolve the QoS configuration challenges of the MMS mission described in Section III-B.

**1. Configurators** determine values for various low-level platform-specific QoS parameters, such as middleware, operating system, and network settings for an application based on its QoS characteristics and requirements, such as relative importance and end-to-end delay. For example, the `MiddlewareConfigurator` configures component Lightweight CCM policies, such as threading policy, priority model, and request processing policy based on the class of the application (*important* and *best-effort*). The `OperatingSystemConfigurator` configures operating system parameters, such as the Rate Monotonic Scheduling (RMS)-based [27] or Maximum Urgency First (MUF)-based [28] priorities of the *Component Servers* that host the components. Likewise, the `NetworkConfigurator` configures network parameters, such as `difserv` code-points of the component interconnections. Like other entities of RACE, Configurators are implemented as CCM components, so new configurators can be plugged into RACE by (re)configuring RACE using PICML.

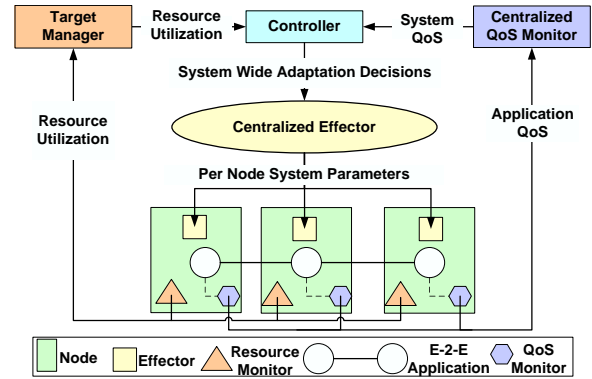
**2. Orchestrator and Conductor.** Based on the QoS properties of the application captured in the E-2-E IDL structure, the Orchestrator selects appropriate Configurators to configure QoS properties for the application. As before, this orchestration is captured in the `Composition IDL` structure and passed onto the Conductor, which invokes the Configurators specified in the `Composition` to configure the system QoS parameters for the application.

RACE’s configurators, orchestrator and conductor coordinate with one another to configure platform-specific QoS parameters for applications appropriately. These entities provide higher level abstractions and shield system developers and SA-POP from low-level platform-specific details, thus resolving the challenges associated with configuring platform-specific QoS parameters identified in Section III-B.

### C. Runtime System Management

When resources are allocated to components at design-time by system designers using PICML, these operations are performed based on estimated resource utilization of applications and estimated availability of system resources. Allocation algorithms supported by RACE’s Allocators allocate resources to components based on current system resource utilization and component’s estimated resource requirements. In open DRE systems, however, there is often no accurate *a priori* knowledge of input workload or the relationship between the resource requirement and QoS of components that make up the application. In these systems, moreover, operational conditions and resource availability cannot be characterized accurately *a priori*.

To resolve the above described challenges, as well as the ones described in III-B, RACE’s control architecture employs a feedback loop to manage system resource and application QoS and ensures (1) QoS requirements of applications are met at all times and (2) system stability by maintaining utilization of system resources below their specified set-points. RACE’s control architecture features a feedback loop that consists of three main components: Monitors, Controllers, and Effectors, as shown in Figure 3.



3: RACE’s Feedback Control Loop

**1. Monitors.** To ensure system stability and meet QoS requirements of applications, RACE’s control architecture must monitor both system QoS and resource utilization. As shown in Figure 3, RACE employs the Lightweight CCM’s `TargetManager` to monitor system resource utilization.

As described in Sidebar 1, containers provide application components with an execution environment and enables them to communicate via the underlying middleware. Each container is aware of all the interactions of a component and the end-to-end delay of an application can therefore be measured in an application-transparent way. QoS properties, such as accuracy, precision, and fidelity of the produced output, are application-specific, however, and thus cannot be measured by the middleware without help from application components. We extended the container to embed Monitors, called *application-QoS-monitors*, to measure end-to-end application delay.

Since QoS-enabled CCM middleware currently implement inter-component interactions (both *facet/receptacle* interactions and *event source/sink* interactions) as two-way calls, end-to-end delay of an application can be obtained by measuring the round-trip delay at the “source” of the application. *Application-QoS-monitors* use high resolution timers (`ACE.High_Res_Timer`) to measure this round-trip delay and periodically send the collected end-to-end delays to the *node-QoS-monitor* that is collocated on the same node using the `Node_QoS_Monitor` interface shown in Figure 4. *Node-QoS-monitors* in turn periodically send the collected end-to-end delay of all the applications on its node to the *centralized-QoS-monitor* using the `Centralized_QoS_Monitor` interface shown in Figure 4. Moreover, application specific QoS monitors can send QoS information to the central monitor by invoking the same interface. The update period of both *application-QoS-monitors* and *node-QoS-monitors* is configurable.

As shown in Figure 3, RACE’s QoS Monitors are structured in the following hierarchical fashion: an *application-QoS-monitor* tracks the QoS of an application, a *node-QoS-monitor* tracks the QoS of all the applications running on its node, and the *centralized-QoS-monitor* tracks the QoS of all the applications running the entire domain, which captures the system QoS. RACE’s Controller(s) obtain the system QoS from the *centralized-QoS-monitor* via the

```

module RACE {
  interface Node_QoS_Monitor {
    // Update period.
    attribute long interval;

    // Oneway method call to push the collected QoS.
    oneway void push_QoS (in string QoS_id,
                        in string Application_id,
                        in any QoS);
  };

  interface Central_QoS_Monitor {
    // No QoS information is available for
    // the requested ApplicationID.
    exception ApplicationIdNotFound { };

    // No QoS information is available for
    // the requested QoSID.
    exception QoSIdNotFound { };

    // Oneway method to push the collected QoS.
    oneway void push_QoS (in string QoS_id,
                        in string Application_id,
                        in any QoS);

    // Retrieve the QoS information regarding a
    // specific QoS of an application.
    any get_QoS (in string QoS_id,
                in string Application_id)
    raises (ApplicationIdNotFound, QoSIdNotFound);
  };
};

```

#### 4: Interface Definition of *Node-QoS-Monitor* and *Centralized-QoS-Monitor*

Central\_QoS\_Monitor interface shown in Figure 4.

**2. Controllers** enable a DRE system to adapt to changing operational context and variations in resource availability and/or demand. The RACE Controllers implement various control algorithms that manage runtime system performance, including EUCON [23], HySUCON [24], and FMUF [29]. Based on the control algorithm they implement, Controllers modify configurable system parameters (such as execution rates and mode of operation of the application), real-time configuration settings (such as operating system priorities of *component servers* that host the components), and network difserv code-points of the component interconnections. Controllers are also implemented as CCM components. RACE can therefore be configured with new Controllers by using PICML.

**3. Effectors** modify system parameters, including resources allocated to components, execution rates of applications, and OS/middleware/network QoS setting for components, to achieve the controller recommended adaptation. As shown in Figure 3, Effectors are designed hierarchically. The *centralized effector* first computes the values of various system parameters for all the nodes in the domain to achieve the Controller recommended adaptation. The computed values of system parameters for each node are then propagated to Effectors located on each node, which then modify system parameters of its node accordingly. The hierarchical design of ResourceMonitors (TargetManager), QoS-Monitors, and RACE’s Effectors is scalable and can handle many applications and nodes in the domain.

**4. Historian, Orchestrator, and Conductor.** The Historian maintains the history of all deployed applications along with their QoS characteristics and mapping of components to nodes. The orchestrator employs the Historian to obtain information regarding the QoS characteristics of application that have been deployed in the system to select the appropriate controller to manage the system. For example, if all the deployed applications can be operated at various rates, the Orchestrator selects the EUCON controller to manage the system. The Conductor invokes the controller selected by the Orchestrator to manage the DRE system.

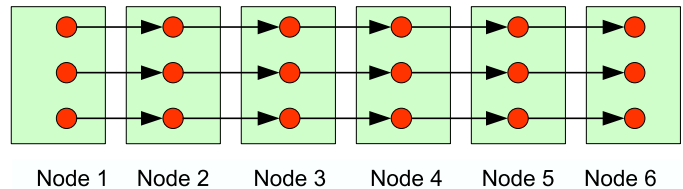
RACE’s monitoring framework, controllers, and effectors coordinate with one another and other entities of RACE to ensure (1) QoS requirements of applications are met and (2) utilization of system resources are maintained within the specified utilization set-point set-point(s), thereby resolving the challenges associated with runtime end-to-end QoS management identified in Section III-B.

## V. PERFORMANCE RESULTS AND ANALYSIS

This section presents the testbed and experiment configuration inspired by the goals of the NASA MMS mission prototype that we developed to evaluate the empirical performance RACE. We describe our experiments and analyze the results to show the performance of this DRE system with and without RACE under varying operating condition and input workload. The results show that RACE performs effective end-to-end adaptation and yields a predictable and high-performance DRE system.

### A. Hardware Testbed

Our experiments were performed on the ISISLab testbed at Vanderbilt University ([www.dre.vanderbilt.edu/ISISlab](http://www.dre.vanderbilt.edu/ISISlab)). The hardware configuration consists of six nodes acting as the system domain. The hardware configuration of all the nodes was a 2.8 GHz Intel Xeon dual processor, 1 GB physical memory, 1GHz Ethernet network interface, and 40 GB hard drive. Redhat Fedora Core release 4 operating system along with real-time patches running in single processor mode was used for all the nodes.



5: Experiment Topology

### B. Summary of Evaluated Scheduling Algorithms

We studied the performance of the prototype MMS system under various configurations: including (1) a baseline configuration with no RACE usage at all and with (2) RACE’s Rate Monotonic Scheduling (RMS) [27] configurator, (3) RACE’s Maximum Urgency First (MUF) [28] configurator, and (4)

RACE’s MUF configurator and Flexible MUF (FMUF) [29] controller. The RMS and the MUF configurators assign priorities to all components (*component servers*) at deployment time based on the RMS and MUF policies, respectively. A disadvantage of RMS is that it cannot provide performance isolation for important applications [28]. During system overload caused by dynamic workload, a important applications with a low rate may miss deadlines, while a best-effort applications with a high rate may experience no missed deadlines.

In contrast, MUF provides performance isolation to important applications by dividing priorities into two classes [28]. All components belonging to important applications are assigned to the high-priority class, while all components belonging to best-effort applications are assigned to the low-priority class. Components within a same priority class are assigned priorities based on the RMS policy. Relative to RMS, however, MUF may cause priority inversion when an important application has a lower rate than best-effort applications. As a result, MUF may unnecessarily cause a best-effort application to miss its deadline, even when all tasks are schedulable under RMS.

To address the limitation of MUF, RACE’s FMUF controller provides performance isolation for important applications while reducing the deadline misses of best-effort applications. While both RMS and MUF assign priorities statically at deployment time, the FMUF controller adjusts the priorities of best-effort applications dynamically based on performance feedback. The FMUF controller can reassign best-effort applications to the high-priority class when (1) all the applications currently in the high-priority class meet their deadlines while (2) some applications in the low-priority class miss their deadlines. FMUF moves best-effort applications back to the low-priority class when the high-priority class experiences deadline misses. It can therefore effectively deal with workload variations caused by application arrivals and changes in application execution times.

### C. Experiment Configuration

Our configurations of the prototype NASA MMS Mission DRE system consist of 11 periodic applications, 4 belonging to important (I-M) class and 7 belonging to best-effort (B-E) class. Each application is composed of 6 components (C1–C6) and is subjected to an end-to-end deadline equal to its period. Interconnection of components of applications is shown in Figure 5. Periods of applications along with the estimated execution times of components comprising the applications are described in Table I. For all applications, static allocation as shown in Figure 5 was specified .

Since the applications described above do not support rate adaptation, RACE employs FMUF as the end-to-end adaptation strategy. Since RACE is a framework, however, other adaptation strategies/algorithms, such as HySUCON [24], can be implemented and employed in a similar way. Below, we evaluate the use of FMUF for end-to-end adaptation. Since the focus of this work is RACE, and not the design or evaluation of individual control algorithms, we use FMUF as an example to demonstrate RACE’s ability to support the integration of

#	Estimated Execution Time (msec)						Period (msec)	Class
	C1	C2	C3	C4	C5	C6		
1	55	40	65	55	40	65	700	I-M
2	90	65	70	90	65	70	1000	I-M
3	65	70	65	70	55	65	900	I-M
4	35	40	40	35	35	40	500	B-E
5	70	65	65	55	65	70	800	B-E
6	70	65	90	70	90	70	1200	B-E
7	40	55	35	40	40	65	600	B-E
8	65	55	55	70	40	55	700	B-E
9	70	65	65	90	70	65	900	B-E
10	35	40	35	35	40	35	400	B-E
11	65	55	65	70	65	70	700	I-M

I: Properties of End-to-End Applications

feedback control algorithms for end-to-end adaptation in DRE systems.

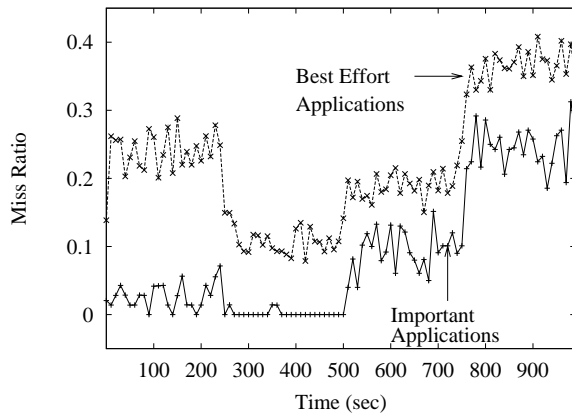
Our experiments were conducted over 1,000 seconds and we emulated the variation in operating condition and input workload by performing the following. At time  $T = 0sec$ , we deployed applications 1 through 10. At time  $T = 250sec$ , we decreased the execution time of all application components by 10 percent to emulate a decrease in input workload. At time  $T = 500sec$ , we deployed application 11, and at time  $T = 750sec$ , we increased the execution time of all application components by 10 percent to emulate an increase in input workload. Since each application was subjected to an end-to-end deadline equal to its period, to evaluate the performance of RACE, we monitored the *deadline miss ratio* of all applications that were deployed.

As described in Sections I and II, QoS-enabled component middleware previously lacked the ability to automatically (1) configure QoS settings of application components and (2) enforce their end-to-end QoS requirements. When the MMS DRE system described above was built atop CIAO/DAnCE directly without RACE, therefore, all application components were assigned default QoS settings, *i.e.* all applications were assigned the same middleware, operating system, and network policies and/or priorities. We use this configuration as the baseline to compare with performance of the system when built atop RACE.

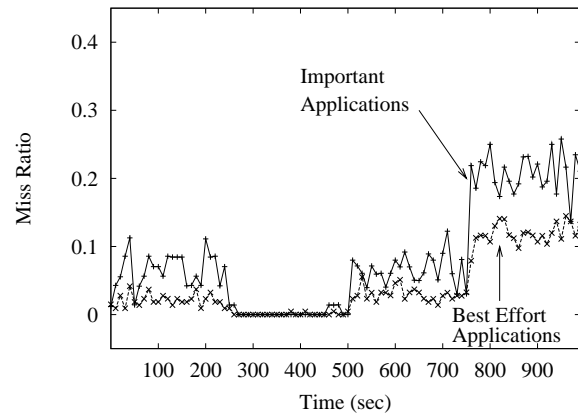
### D. Analysis of Empirical Results

We now present the results obtained from running the experiment described in Section V-C on our ISISlab DRE system testbed described in Section V-A. We use deadline miss ratio as the metric to evaluate system performance under varying input workloads and operating conditions.

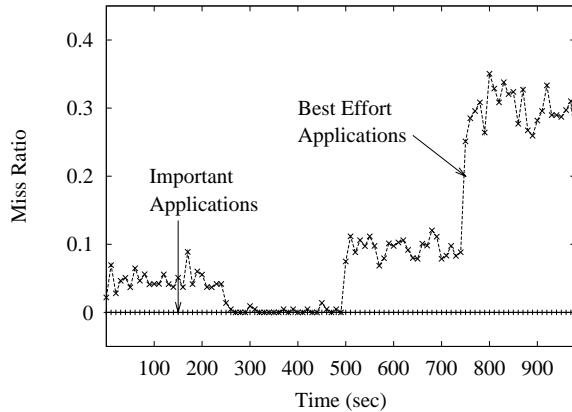
**Comparison of QoS.** Figures 6a, 6b, 6c, and 6d show the deadline miss ratio of applications when the system operated under the four configurations described in Section V-C, *i.e.*, baseline configuration, with RACE’s RMS configurator, RACE’s MUF configurator, and RACE’s MUF configurator along with FMUF controller, respectively. These figures show that under all the four configurations, deadline miss ratio of applications (1) reduced at  $T = 250sec$  due to the decrease in the input work load, (2) increased at  $T = 500sec$  due to the introduction of application 11, and (3) further increased at  $T = 750sec$  due to the increase in the input workload.



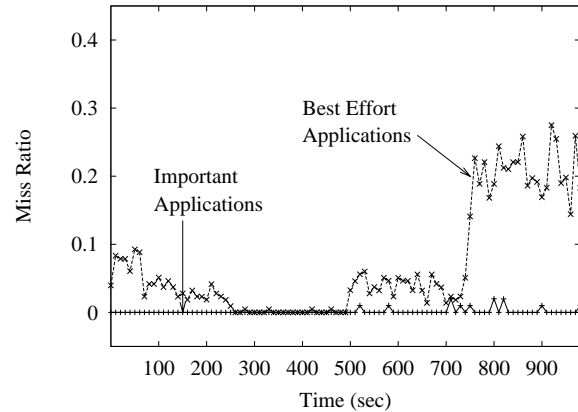
(a) Baseline



(b) RMS Configurator



(c) MUF Configurator



(d) MUF Configurator + FMUF Controller

## 6: Deadline Miss Ratio

These results demonstrate the impact of fluctuation in input workload and operating conditions on system performance.

Figure 6b shows that when RACE's RMS configurator was used to configure the operating system priorities of component servers, deadline miss ratio of important applications was higher than that of best-effort applications due to reasons explained earlier. Figures 6c and 6d show that when RACE's MUF configurator is used (both individually and along with FMUF controller), deadline miss ratio of important applications were nearly zero throughout the course of the experiment. Figures 6a and 6c demonstrate that RACE improves QoS of our DRE system significantly by configuring platform-specific parameters appropriately.

As described in [29], the FMUF controller responds to variations in input workload and operating conditions (indicated by deadline misses) by dynamically adjusting the priorities of the best-effort applications (*i.e.* moving best-effort applications into or out of the high-priority class). Figures 6a and 6d demonstrate the impact of the RACE's controller on the performance of the system.

Our conclusion from analyzing the results above is that RACE significantly improves the performance of our prototype MMS DRE system even under varying input workload and operating conditions. These benefits result from configuring platform-specific QoS parameters appropriately and performing effective end-to-end adaptation, which were carried out by RACE's MUF Configurator and FMUF Controller

respectively. Moreover, the RACE framework addresses the challenges of building component-based DRE systems identified in Section III-B.

### E. Overhead of the RACE Framework.

The runtime overhead of the RACE framework can be decomposed into *monitoring overhead* (the average increase in end-to-end delay as a result of RACE's monitoring framework, as perceived by the application) and *control overhead* (the average execution time of RACE's controller). To measure monitoring overhead, we first instrumented application components with high resolution timers (`ACE_High_Res_Timer`) to measure the end-to-end delay. We next obtained the average end-to-end delay of applications described in Section V-C when the system was executed with and without RACE's monitoring framework. We then computed RACE's monitoring overhead as the difference between these two average end-to-end execution times. To measure the control overhead, we also instrumented RACE's controller with a high resolution timer (`ACE_High_Res_Timer`). This timer measured the execution time of the controller during every sampling period. We computed the control overhead as the average of the collected execution times.

From running the experiment described in Section V-C on our DRE system testbed, the average monitoring overhead was  $37.97 \mu s$  and average control overhead of the FMUF controller was  $799.82 ns$ . Average monitoring and control

overhead are 0.0645% and 0.0013% of the average estimated execution times of components shown in Table I. These results demonstrate that the runtime overhead of RACE is small and acceptable for DRE systems.

## VI. CONCLUDING REMARKS

In this paper, we described RACE, which is an adaptive resource management framework that provides end-to-end adaptation and resource management for open DRE systems built atop QoS-enabled component middleware. We demonstrated how RACE helps resolve key resource and QoS management challenges associated with a prototype of the NASA MMS system. We also discussed results from empirical studies of the overhead associated with RACE.

Since the elements of RACE are designed and implemented as CCM components, RACE itself can be configured using DSML tools, such as PICML. Moreover, new InputAdapters, Allocators, Configurators, and Controllers can be plugged into RACE using PICML, and without any modifications to the existing architecture. RACE can be used to deploy, allocate resources to, and manage performance of, applications that are composed both at design time as well as at runtime. Moreover, due to the ease with which RACE can be configured, RACE can be employed in a wide range of DRE systems.

Our experience building a prototype of a representative DRE system atop RACE shows that it yields in a predictable and high performance system, even in the face of changing operational conditions, workloads, and resource availability. CIAO, DANCE, and RACE are available in open-source for download at [deuce.doc.wustl.edu/Download.html](http://deuce.doc.wustl.edu/Download.html).

## REFERENCES

- [1] *Real-time CORBA Specification*, OMG Document formal/05-01-04 ed., Object Management Group, Aug. 2002.
- [2] G. Bollella, J. Gosling, B. Brosgol, P. Dibble, S. Furr, D. Hardin, and M. Turnbull, *The Real-time Specification for Java*. Addison-Wesley, 2000.
- [3] *Light Weight CORBA Component Model Revised Submission*, OMG Document realtime/03-05-05 ed., Object Management Group, May 2003.
- [4] D. C. Sharp, E. Pla, K. R. Luecke, and R. J. H. II, "Evaluating Real-time Java for Mission-Critical Large-Scale Embedded Systems," in *IEEE Real-time and Embedded Technology and Applications Symposium*. Washington, DC: IEEE Computer Society, May 2003.
- [5] D. C. Sharp and W. C. Roll, "Model-Based Integration of Reusable Component-Based Avionics System," in *Proc. of the Workshop on Model-Driven Embedded Systems in RTAS 2003*, May 2003.
- [6] 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 *Proceedings of the 11th Real-time Technology and Application Symposium (RTAS '05)*. San Francisco, CA: IEEE, Mar. 2005, pp. 190–199.
- [7] N. Wang, D. C. Schmidt, A. Gokhale, C. Rodrigues, B. Natarajan, J. P. Loyal, R. E. Schantz, and C. D. Gill, "QoS-enabled Middleware," in *Middleware for Communications*, Q. Mahmoud, Ed. New York: Wiley and Sons, 2004, pp. 131–162.
- [8] J. Kinnebrew, N. Shankaran, G. Biswas, and D. Schmidt, "A Decision-Theoretic Planner with Dynamic Component Reconfiguration for Distributed Real-Time Applications," in *Poster paper at the Twenty-First National Conference on Artificial Intelligence*, Boston, MA, July 2006.
- [9] *Deployment and Configuration Adopted Submission*, Document ptc/03-07-08 ed., OMG, July 2003.
- [10] V. F. Wolfe, L. C. DiPippo, R. Bethmagalkar, G. Cooper, R. Johnston, P. Kortmann, B. Watson, and S. Wohlever, "RapidSched: Static Scheduling and Analysis for Real-Time CORBA," *WORDS*, vol. 00, p. 34, 1999.
- [11] J. W. Liu, J. Redondo, Z. Deng, T. Tia, R. Bettati, A. Silberman, M. Storch, R. Ha, and W. Shih, "PERTS: A Prototyping Environment for Real-Time Systems," Campaign, IL, USA, Tech. Rep., 1993.
- [12] B. Ravindran, L. Welch, and B. Shirazi, "Resource Management Middleware for Dynamic, Dependable Real-Time Systems," *Real-Time Syst.*, vol. 20, no. 2, pp. 183–196, 2001.
- [13] L. R. Welch, B. A. Shirazi, B. Ravindran, and C. Bruggeman, "DeSiDeRaTa: QoS Management Technology for Dynamic, Scalable, Dependable Real-time Systems," in *IFACs 15th Symposium on Distributed Computer Control Systems (DCCS98)*. IFAC, 1998.
- [14] D. Fleeman, M. Gillen, A. Lenharth, M. Delaney, L. R. Welch, D. Juedes, and C. Liu, "Quality-Based Adaptive Resource Management Architecture (QARMA): A CORBA Resource Management Service," *IPDPS*, vol. 03, p. 116b, 2004.
- [15] C. D. Gill, "Flexible Scheduling in Middleware for Distributed Rate-Based Real-time Applications," Ph.D. dissertation, Department of Computer Science, Washington University, St. Louis, 2002.
- [16] K. Bryan, L. C. DiPippo, V. Fay-Wolfe, M. Murphy, J. Zhang, D. Niehaus, D. T. Fleeman, D. W. Juedes, C. Liu, L. R. Welch, and C. D. Gill, "Integrated CORBA Scheduling and Resource Management for Distributed Real-Time Embedded Systems," in *RTAS '05: Proceedings of the 11th IEEE Real Time on Embedded Technology and Applications Symposium*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 375–384.
- [17] 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*, Portland, OR, May 2003.
- [18] 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 *Proceedings of the IEEE Real-time Applications Symposium*. Washington, DC: IEEE, May 2003, pp. 58–69.
- [19] 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)*. Washington, DC: IEEE, May 2003.
- [20] R. Schantz, J. Loyal, M. Atighetchi, and P. Pal, "Packaging Quality of Service Control Behaviors for Reuse," in *Proceedings of the 5th IEEE International Symposium on Object-Oriented Real-time Distributed Computing (ISORC)*. Crystal City, VA: IEEE/IFIP, April/May 2002, pp. 375–385.
- [21] X. Wang, C. Lu, and X. Koutsoukos, "Enhancing the Robustness of Distributed Real-Time Middleware via End-to-End Utilization Control," in *RTSS '05: Proceedings of the 26th IEEE International Real-Time Systems Symposium*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 189–199.
- [22] N. Shankaran, X. Koutsoukos, C. Lu, D. C. Schmidt, and Y. Xue, "Hierarchical Control of Multiple Resources in Distributed Real-time and Embedded Systems," in *Proceedings of the Euromicro Conference on Real-Time Systems (ECRTS 06)*, Dresden, Germany, July 2006.
- [23] C. Lu, X. Wang, and X. Koutsoukos, "Feedback Utilization Control in Distributed Real-time Systems with End-to-End Tasks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 16, no. 6, pp. 550–561, 2005.
- [24] X. Koutsoukos, R. Tekumalla, B. Natarajan, and C. Lu, "Hybrid Supervisory Control of Real-time Systems," in *11th IEEE Real-time and Embedded Technology and Applications Symposium*, San Francisco, California, Mar. 2005.
- [25] *Deployment and Configuration Adopted Submission*, OMG Document mars/03-05-08 ed., Object Management Group, July 2003.
- [26] D. de Niz and R. Rajkumar, "Partitioning Bin-Packing Algorithms for Distributed Real-time Systems," *International Journal of Embedded Systems*, 2005.
- [27] J. Lehoczky, L. Sha, and Y. Ding, "The Rate Monotonic Scheduling Algorithm: Exact Characterization and Average Case Behavior," in *Proceedings of the 10th IEEE Real-time Systems Symposium (RTSS 1989)*. IEEE Computer Society Press, 1989, pp. 166–171.
- [28] D. B. Stewart and P. K. Khosla, "Real-time Scheduling of Sensor-Based Control Systems," in *Real-time Programming*, W. Halang and K. Ramamritham, Eds. Tarrytown, NY: Pergamon Press, 1992.
- [29] Y. Chen and C. Lu, "Flexible Maximum Urgency First Scheduling for Distributed Real-Time Systems," Washington University in St. Louis, Tech. Rep. WUCSE-2006-55, October 2006.