

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCSE-2008-5

2008-01-01

Reconfigurable Real-Time Middleware for Distributed Cyber-Physical Systems with Aperiodic Events

Yuanfang Zhang, Christopher Gill, and Chenyang Lu

Different distributed cyber-physical systems must handle aperiodic and periodic events with diverse requirements. While existing real-time middleware such as Real-Time CORBA has shown promise as a platform for distributed systems with time constraints, it lacks flexible configuration mechanisms needed to manage end-to-end timing easily for a wide range of different cyber-physical systems with both aperiodic and periodic events. The primary contribution of this work is the design, implementation and performance evaluation of the first configurable component middleware services for admission control and load balancing of aperiodic and periodic event handling in distributed cyber-physical systems. Empirical results demonstrate the need... **Read complete abstract on page 2.**

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research



Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Zhang, Yuanfang; Gill, Christopher; and Lu, Chenyang, "Reconfigurable Real-Time Middleware for Distributed Cyber-Physical Systems with Aperiodic Events" Report Number: WUCSE-2008-5 (2008). *All Computer Science and Engineering Research*.
https://openscholarship.wustl.edu/cse_research/236

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

Reconfigurable Real-Time Middleware for Distributed Cyber-Physical Systems with Aperiodic Events

Yuanfang Zhang, Christopher Gill, and Chenyang Lu

Complete Abstract:

Different distributed cyber-physical systems must handle aperiodic and periodic events with diverse requirements. While existing real-time middleware such as Real-Time CORBA has shown promise as a platform for distributed systems with time constraints, it lacks flexible configuration mechanisms needed to manage end-to-end timing easily for a wide range of different cyber-physical systems with both aperiodic and periodic events. The primary contribution of this work is the design, implementation and performance evaluation of the first configurable component middleware services for admission control and load balancing of aperiodic and periodic event handling in distributed cyber-physical systems. Empirical results demonstrate the need for, and the effectiveness of, our configurable component middleware approach in supporting different applications with aperiodic and periodic events, and providing a flexible software platform for distributed cyber-physical systems with end-to-end timing constraints.

2008-5

Reconfigurable Real-Time Middleware for Distributed Cyber-Physical Systems with Aperiodic Events

Authors: Yuanfang Zhang, Christopher Gill and Chenyang Lu

Corresponding Author: {yfzhang, cdgill, lu}@cse.wustl.edu

Abstract: Different distributed cyber-physical systems must handle aperiodic and periodic events with diverse requirements. While existing real-time middleware such as Real-Time CORBA has shown promise as a platform for distributed systems with time constraints, it lacks flexible configuration mechanisms needed to manage end-to-end timing easily for a wide range of different cyber-physical systems with both aperiodic and periodic events. The primary contribution of this work is the design, implementation and performance evaluation of the first configurable component middleware services for admission control and load balancing of aperiodic and periodic event handling in distributed cyber-physical systems. Empirical results demonstrate the need for, and the effectiveness of, our configurable component middleware approach in supporting different applications with aperiodic and periodic events, and providing a flexible software platform for distributed cyber-physical systems with end-to-end timing constraints.

Type of Report: Other

Reconfigurable Real-Time Middleware for Distributed Cyber-Physical Systems with Aperiodic Events *

Yuanfang Zhang, Christopher Gill and Chenyang Lu
Department of Computer Science and Engineering
Washington University, St. Louis, MO, USA
{yfzhang, cdgill, lu}@cse.wustl.edu

Abstract

Different distributed cyber-physical systems must handle aperiodic and periodic events with diverse requirements. While existing real-time middleware such as Real-Time CORBA has shown promise as a platform for distributed systems with time constraints, it lacks flexible configuration mechanisms needed to manage end-to-end timing easily for a wide range of different cyber-physical systems with both aperiodic and periodic events. The primary contribution of this work is the design, implementation and performance evaluation of the first configurable component middleware services for admission control and load balancing of aperiodic and periodic event handling in distributed cyber-physical systems. Empirical results demonstrate the need for, and the effectiveness of, our configurable component middleware approach in supporting different applications with aperiodic and periodic events, and providing a flexible software platform for distributed cyber-physical systems with end-to-end timing constraints.

1 Introduction

Many distributed cyber-physical systems (CPS) must handle a mix of aperiodic and periodic events, including aperiodic events with end-to-end deadlines whose assurance is critical to the correct behavior of the system. Computer-integrated manufacturing is representative of many distributed CPS. For example, in an industrial plant monitoring system, an aperiodic alert may be generated when a series of periodic sensor readings meets certain hazard detection criteria. This alert must be processed on multiple processors within an end-to-end deadline, e.g., to put an industrial process into a fail-safe mode. User inputs and other sensor readings may trigger other real-time aperiodic events.

While traditional real-time middleware such as Real-Time CORBA [20] and Real-Time Java [5] have shown promise as distributed software platforms for systems with time constraints, existing middleware systems lack the flexibility needed to support cyber-physical systems with diverse application semantics and requirements. For example, load balancing is an effective mechanism for handling variable real-time workloads in a distributed cyber-physical system. However, its suitability for cyber-physical systems highly depends on their application semantics. Some digital control algorithms (e.g., proportional-integral-derivative control) for physical systems are stateful and hence not amenable for frequent task re-allocation caused by load balancing, while others (e.g., proportional control) do not have such limitations. Similarly, job skipping (skipping the processing of certain instances of a periodic task) is an admission control strategy for dealing with transient system overload. However, job skipping is not suitable for certain critical control applications in which missing one job may cause catastrophic consequences on the controlled system. In contrast, other applications ranging from video reception to telecommunications may be able to tolerate varying degrees of job skipping [17].

Therefore, a key open challenge for distributed cyber-physical systems is to develop a flexible middleware infrastructure that can be easily configured to support the diverse requirements of cyber-physical systems. Specifically, middleware services such as load balancing and admission control must support a variety of strategies. Furthermore, the configuration of those strategies must be supported in a flexible yet principled way, so that system developers are able to explore alternative configurations but invalid configurations cannot be chosen by mistake.

Providing middleware services with configurable strategies thus faces several important challenges: (1) services must be able to provide configurable strategies, and configuration tools must be added or extended to allow configuration of those strategies; (2) the specific criteria that distinguish which service strategies are preferable must be

*This work was supported in part by NSF grant CCF-0615341 and NSF CAREER award CNS-0448554.

identified, and applications must be categorized according to those criteria; and (3) appropriate combinations of services' strategies must be identified for each such application category, according to its characteristic criteria. To address these challenges, and thus to enhance support for diverse cyber-physical systems with aperiodic and periodic events, we have designed and implemented a new set of component middleware services including end-to-end event scheduling, admission control, and load balancing. We have also developed configuration tools to integrate these service components for each particular application according to its specific criteria.

Research Contributions: In this work, we have (1) developed what is to our knowledge the first set of configurable component middleware services supporting multiple admission control and load balancing strategies for handling aperiodic and periodic events; (2) developed a novel component configuration pre-parser and interfaces to configure real-time admission control and load balancing services flexibly at system deployment time; (3) defined categories of cyber-physical applications according to specific characteristics, and related them to suitable combinations of strategies for our services; and (4) provided a case study that applies different configurable services to a domain with both aperiodic and periodic events, offers empirical evidence of the overheads involved and the trade-offs among service configurations, and demonstrates the effectiveness of our approach in that domain. Our work thus significantly enhances the applicability of real-time middleware as a flexible infrastructure for distributed cyber-physical systems.

Section 2 introduces the middleware systems and scheduling theory underlying our middleware services. Section 3, 4 and 5 present our middleware architecture, configurable strategies, and component implementations for supporting end-to-end event handling in cyber-physical systems. Section 6 describes our new configuration engine extensions, which can flexibly configure different strategies for our services according to each application's requirements. Section 7 evaluates the performance of our approach, including trade-offs among different service strategy combinations, and characterizes the overheads introduced by our approach. Section 8 presents a survey of related work, and we offer concluding remarks in Section 9.

2 Background

Task Model: We consider cyber-physical systems comprised of physical systems generating aperiodic and periodic events that must be processed on distributed computing platforms subject to end-to-end timing constraints. Henceforth the processing of a sequence of events is referred to as a *task*. A task T_i is composed of a chain of subtasks $T_{i,j}$ ($1 \leq j \leq n_i$) located on different processors. The

release for execution of a subtask T_i is triggered by the completion of its predecessor $T_{i,j-1}$, where processing one event on the sequence implements a subtask. Each release of a subtask is called one subjob, and each release of a task is a job composed of a chain of subjobs. A task is subject to an end-to-end deadline that is its maximum allowable response time. A subtask has a specified execution time, which is the maximum execution of all the subjobs in it. The period of a periodic task is the interarrival time of consecutive subjobs of the first subtask in the periodic task. Aperiodic tasks do not have periods. Their interarrival times between consecutive subjobs of the first subtask may vary widely and, in particular, can be arbitrary small. The execution time of every subtask, the end-to-end deadline of every task, and the period of every periodic task in the system are known.

Component Middleware: Component middleware platforms are an effective way of achieving customizable reuse of software artifacts. In these platforms, *components* are units of implementation and composition that collaborate with other components via *ports*. The ports isolate the components' contexts from their actual implementations. Component middleware platforms provide execution environments and common services, and support additional tools to configure and deploy the components.

In previous work we developed the first instantiation of a middleware admission control service supporting both aperiodic and periodic events [30] (on TAO, a widely used Real-Time CORBA middleware). However, our previous admission control service only included a fixed set of strategies. As is shown in Section 4, a diverse set of inter-operating services and strategies is needed to support cyber-physical systems with different application semantics. Unfortunately, it is difficult to extend implementations that rely directly on distributed object middleware such as our original admission control service. Specifically, in those middleware systems changing the supported strategy requires explicit changes to the service code itself, which can be tedious and error-prone in practice.

The Component-Integrated ACE ORB (CIAO) [14] implements the Light Weight CCM specification [21] and is built atop the TAO [15] real-time CORBA object request broker (ORB). CIAO abstracts real-time policies as installable and configurable units. However, CIAO does not support aperiodic task scheduling, admission control or load balancing. To develop a flexible infrastructure for cyber-physical systems, in this work we develop new admission control and load balancing services, each with a set of alternative strategies on top of CIAO. Furthermore, we extended CIAO to configure and manage both services.

DAnCE [9] is a QoS-enabled component deployment and configuration engine that implements the Object Management Group (OMG)'s Light Weight Deployment and

Configuration specification [21]. DAnCE parses component configuration/deployment descriptions and automatically configures and deploys ORBs, containers, and server resources at system initialization time, to enforce end-to-end QoS requirements. However, DAnCE does not provide essential features needed to configure our admission control and load balancing services correctly, e.g., to disallow invalid combinations of service strategies.

Aperiodic Task Support: Aperiodic tasks have been studied extensively in real-time scheduling theory, including work on aperiodic servers that integrate scheduling of aperiodic and periodic tasks [24]. New schedulability tests based on aperiodic utilization bounds [1] and a new admission control approach [4] also were introduced recently. In our previous work [30], we implemented and evaluated an admission control services for two suitable aperiodic scheduling techniques (aperiodic utilization bound [1] and deferrable server [25]) on TAO. Since aperiodic utilization bound (AUB) has a comparable performance to deferrable server, and requires less complex scheduling mechanisms in middleware, we focus exclusively on the AUB scheduling technique in this paper. Our experiences with AUB reported in this paper show how configurability of other techniques can be integrated within real-time component middleware in a similar way.

With the AUB approach, three kinds of strategies must be made configurable to provide flexible and principled support for diverse cyber-physical systems with aperiodic and periodic tasks: (1) when admissibility is evaluated (to trade-off the granularity and thus the pessimism of admission guarantees), (2) when the contributions of completed jobs of subtasks can be removed from the schedulability analysis used for admission control (to reduce pessimism), and (3) when tasks can be assigned to different processors (to balance load and improve system performance).

In AUB [1], the set of *current* tasks $S(t)$ at any time t is defined as the set of tasks that have been released but whose deadlines have not expired. Hence, $S(t) = \{T_i | A_i \leq t < A_i + D_i\}$, where A_i is the release time of the first subtask of task T_i , and D_i is the deadline of task T_i . The synthetic utilization of processor j at time t , $U_j(t)$, is defined as the sum of individual subtask utilizations on the processor, accrued over all current tasks. According to AUB analysis, a system achieves its highest schedulable synthetic utilization bound under End-to-end Deadline Monotonic Scheduling (EDMS) algorithm under certain assumptions. Under EDMS, a subtask has a higher priority if it belongs to a task with a shorter end-to-end deadline. Note that AUB does not distinguish aperiodic from periodic tasks. All tasks are scheduled using the same scheduling policy. Under EDMS task T_i will meet

its deadline if the following condition holds [1]:

$$\sum_{j=1}^{n_i} \frac{U_{V_{ij}}(1 - U_{V_{ij}}/2)}{1 - U_{V_{ij}}} \leq 1 \quad (1)$$

where V_{ij} is the j^{th} processor that task T_i visits. A task (or an individual job) can be admitted only when this condition continues to be satisfied for all admitted tasks and this task. Since applications may or may not tolerate job skipping, whether this condition is checked only when a task first arrives or whenever each job arrives should be configurable.

Note that a task remains in the current task set even if it has been completed, as long as its deadline has not expired. To reduce the pessimism of the AUB analysis, a *resetting rule* is introduced in [1]. When a processor becomes idle, the contribution of all completed subtasks to the processor's synthetic utilization can be removed without affecting the correctness of the schedulability condition (1). Since the resetting rule introduces extra overhead, it should be made configurable whether the contribution of only aperiodic subtasks or of both aperiodic and periodic subtasks can be removed early. Under AUB-based schedulability analysis, load balancing also can effectively improve system performance [1]. However some applications require persistent state preservation between jobs of the same task, so it should be made configurable whether a task can be re-allocated to a different processor at each release.

3 Overview of Middleware Architecture

To support end-to-end aperiodic and periodic events in diverse cyber-physical applications, we have developed a new middleware architecture for task management, and a corresponding front-end configuration engine. The key feature of our approach is a *configurable component framework* that can be customized for different sets of aperiodic and periodic tasks. Our framework provides configurable *admission controller* (AC), *idle resetter* (IR) and *load balancer* (LB) components which interact with application components through *task effector* (TE) components. The AC component provides on-line admission control and schedulability tests for tasks that arrive dynamically at run time. The LB component provides an acceptable task assignment plan to the admission controller if the new arrival task is admissible. The IR component reports all completed subtasks on one processor to the AC component when the processor becomes idle, so the AC component can remove their expected utilization to reduce the pessimism of the AUB analysis at run-time.

Figure 1 illustrates our distributed middleware architecture. All processors are connected by TAO's federated event channel [13] which pushes events through local event channels, gateways and remote event channels to the events'

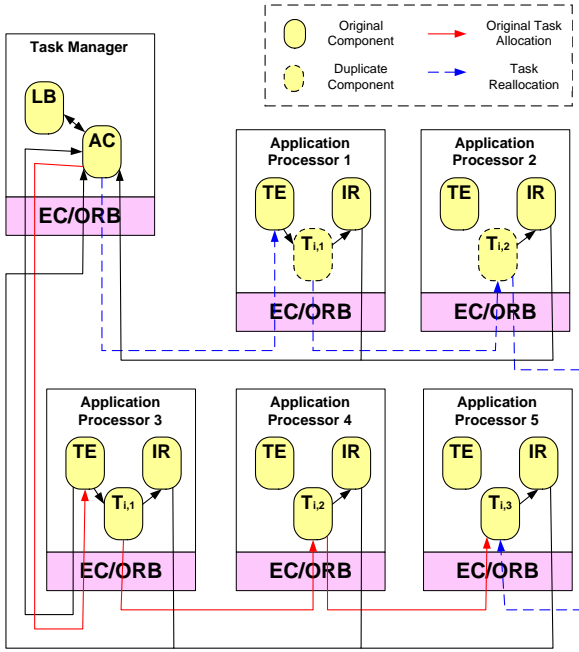


Figure 1. Distributed Middleware Architecture

consumers sitting on different processors. We deploy one AC component and one LB component on a central task manager processor, and one IR component and one task effector (TE) component on each of multiple application processors. Figure 1 shows an example end-to-end task T_i composed of 3 consecutive subtasks, $T_{i,1}$, $T_{i,2}$ and $T_{i,3}$, executing on separate processors. $T_{i,1}$ and $T_{i,2}$ have duplicates on other application processors. When task T_i arrives at an application processor, the task effector component on that processor pushes a “Task Arrive” event to the AC component and holds the task until it receives an “Accept” command from the AC component. The AC component and LB component decide whether to accept the task, and if so, where to assign its subtasks. The solid line and the dashed line show two possible assignments of subtasks. If the first subtask $T_{i,1}$ is not assigned to the processor where T_i arrived, we call this assignment a *task re-allocation*.

An advantage of this centralized AC/LB architecture is that it does not require synchronization among distributed admission controllers. In contrast, in a distributed architecture the AC components on multiple processors may need to coordinate and synchronize with each other in order to make correct decisions, because admitting an end-to-end task may affect the schedulability of other tasks located on the multiple affected processors. A potential disadvantage of the centralized architecture is that the AC component may become a bottleneck and thus affect scalability. However, the computation time of the schedulability analysis is significantly lower than task execution times in many distributed cyber-physical systems, which alleviates the scalability limitations

of a centralized solution [30]. In summary, while our real-time component middleware approach can be extended to use a more distributed architecture, we have adopted a centralized approach with less complexity and overhead, which allows us to focus on achieving system flexibility through component middleware services.

4 Mapping CPS Characteristics to Middleware Strategies

A key contribution of this paper is categorizing a set of characteristics of different cyber-physical systems and mapping them to suitable middleware service strategies. In this section, we present a set of criteria used to categorize cyber-physical systems, and analyze how to map those criteria to different service strategies supported by our middleware.

4.1 CPS Characteristics

Three criteria distinguish how different cyber-physical systems with aperiodic tasks should be supported: Job skipping (criterion C1); State persistency (criterion C2); and Component replication (criterion C3).

Job Skipping means that some jobs of a task are executed while other jobs of the same task may not be admitted. Some applications, such as video streaming, and other loss-tolerant forms of sensing can tolerate job skipping, while in critical control applications, once a task is admitted, all its jobs should be allowed to execute.

State Persistency means that states are required to be preserved between jobs of a same task. For proportional control systems [22], tasks are stateless and only require current information, so jobs can be re-allocated dynamically. However, for integral control systems [22], tasks require incremental calculation and are not suitable for job re-allocation.

Component Replication depends on an application’s throughput requirements. Replication is used here to reduce latency through load distribution, not for fault tolerance purposes. Only those applications with replicated components can support task re-allocation, whereas those that cannot be replicated (e.g. due to constraints on the locality of sensors or actuators) cannot support task re-allocation.

According to these different application criteria, the AC, IR and LB components can be configured to use different strategies. For each component, which strategy is more suitable depends on these criteria and the application’s overhead depends on these criteria and the application’s overhead constraints. Table 1 shows how these criteria help to classify CPS applications, which in turn allows selection of corresponding middleware strategies. We have designed all strategies with corresponding configurable attributes, and provide a configuration pre-parser and a component configuration interface to allow developers to select and configure each service flexibly, according to each application’s

specific needs. We now examine the different strategies for each component and the trade-offs among them.

	No	Yes
C1:Job Skipping	AC per Task	AC per Job
C2:State Persistency	LB per Job	LB per Task
C3:Component Replication	No LB	LB

Table 1. Criteria and Middleware Strategies

4.2 Admission Control (AC) Strategies

Admission control offers significant advantages for systems with aperiodic and periodic tasks, by providing on-line schedulability guarantees to tasks arriving dynamically. Our AC component supports two different strategies: **AC per Task** and **AC per Job**. **AC per Task** performs the admission test only when a task first arrives while **AC per Job** performs the admission test whenever a job of the task arrives. Only applications satisfying criterion C1 are suitable for the second strategy, since it may not admit some jobs. Moreover, the second strategy reduces pessimism at the cost of increasing overhead. The application developer thus needs to consider trade-offs between overhead and pessimism in choosing a proper configuration.

AC per Task: Considering the admission overhead and the fixed inter-arrival times of periodic tasks, one strategy is to perform an admission test only when a periodic task first arrives. Once a periodic task passes the admission test, all its jobs are allowed to be released immediately when they arrive. This strategy improves middleware efficiency at the cost of increasing the pessimism of the admission test. In the AUB analysis [1], the contribution of a job to the synthetic utilization of a processor can be removed when the job’s deadline expires (or when the CPU idles if the resetting rule is used and the subjob has been completed). If admission control is performed only at task arrival time, however, the AC component must reserve the synthetic utilization of the task throughout its lifetime. As a result, it cannot reduce the synthetic utilization between the deadline of a job and the arrival of the subsequent job of the same task, which may result in pessimistic admission decisions [1].

AC per Job: If it is possible to skip a job of a periodic task (criterion C1), the alternative strategy to reduce pessimism is to apply the admission test to every job of a *periodic* task. This strategy is practical for many systems, since the AUB test is highly efficient when used for AC, as is shown in Section 7.3 by our overhead measurements.

4.3 Idle Resetting (IR) Strategies

The use of a resetting rule can reduce the pessimism of the AUB schedulability test significantly [1, 30]. There are

three ways to configure IR components in our approach. The first of these three strategies avoids the resetting overhead, but is the most pessimistic. The third strategy removes the contribution of completed aperiodic and periodic subjobs more frequently than the other two strategies. Although it has the least pessimism, it introduces the most overhead. The second strategy offers a trade-off between the first and the third strategies.

No IR: The first strategy is to use no resetting at all, so that if the subjobs complete their executions, the contributions of completed subjobs to the processor’s synthetic utilization are not removed until the job deadline. This strategy avoids the resetting overhead, but increases the pessimism of schedulability analysis.

IR per Task: The second strategy is that each IR component records completed aperiodic subjobs on one processor. Whenever the processor is idle, a lowest priority thread called an *idle detector* begins to run, and reports the completed *aperiodic* subjobs to the AC component through an “Idle Resetting” event. To avoid reporting repeatedly, the idle detector only reports when there is a newly completed aperiodic subjob whose deadline has not expired.

IR per Job: The third strategy is that each IR component records and reports not only the completed aperiodic subjobs but also the completed subjobs of *periodic* subtasks.

4.4 Load Balancing (LB) Strategies

Under AUB-based AC, load balancing can effectively improve system performance in the face of dynamic task arrivals [1]. We use a heuristic algorithm to assign subtasks to processors at run-time, which always assigns a subtask to the processor with the lowest synthetic utilization among all processors on which the application component corresponding to the task has been replicated (criterion C3).¹ Since migrating a subtask between processors introduces extra overhead, when we accept a new task, we only determine the assignment of that new task and do not change the assignment plan for any other task in the current task set. This service also has three strategies. The first strategy is suitable for applications which cannot satisfy criterion C3. The second strategy is most applicable for applications which only satisfy C3, but can not satisfy criterion C2. The third strategy is most suitable for applications which satisfy both C3 and C2.

No LB: This strategy does not perform load balancing. Each subtask does not have a replica and is assigned to a particular processor.

LB per Task: Each task will only be assigned once, at its first arrival time. This strategy is suitable for applications

¹The focus here is not on the load balancing algorithms themselves. Our configurable middleware may be easily extended to incorporate LB components implementing other load balancing algorithms according to each application’s needs.

which must maintain persistent state between any two consecutive jobs of a periodic task.

LB per Job: The third strategy is the most flexible. All jobs from a periodic task are allowed to be assigned to different processors when they arrive.

4.5 Combining AC, IR and LB Strategies

When we use the AC, IR and LB components together, their strategies can be configured in 18 different combinations. However, some combinations of the strategies are invalid. The AC-per-Task/IR-per-Job combination is not reasonable, because per job idle resetting means the synthetic utilizations of all completed subjobs of periodic subtasks are to be removed from the central admission controller, but per task admission control requires that the admission controller reserves the synthetic utilization for all accepted periodic tasks, so an accepted periodic task does not need to go through admission control again before releasing its jobs. These two requirements are thus contradictory, and we can exclude the corresponding configurations as being invalid. Removing this invalid AC/IR combination means removing 3 invalid AC/IR/LB combinations, so there are only 15 reasonable combinations of strategies left. Looking to this degree of complexity to make “right” performant design decision, an application developer needs definitely some kind of cognitive support. An advantage of our middleware architecture and configuration engine is that they allow application developers to configure middleware services to achieve any valid combination of strategies, while disallowing invalid combinations.

As Figure 2 shows, we divided scheduling of real-time tasks into three axes of configurability: admission control, idle resetting and load balancing. Different configuration options in each of these axes and the impact they may have, as well as conflicting configurations, are delineated thoroughly in this section.

5 Component Implementation

Configurable component middleware standards, such as the CORBA Component Model (CCM) [19], can help to reduce the complexity of developing distributed cyber-physical systems by defining a component-based programming paradigm. They also help by defining a standard configuration framework for packaging and deploying reusable software components. The Component Integrated ACE ORB (CIAO) [28] is an implementation of the Light Weight CCM specification [21] that is suitable for distributed cyber-physical systems. To support the different strategies described in Section 4, and to allow flexible configuration of suitable combinations of those strategies for a variety of applications, we have implemented admission control,




Figure 2. Middleware Services and Strategies

idle resetting and load balancing in CIAO as configurable components. Each component provides a specific service with configurable attributes and clearly defined interfaces for collaboration with other components, and can be instantiated multiple times with the same or different attributes. Component instances can be connected together at runtime through appropriate ports to form a distributed cyber-physical system.

As Figure 3 illustrates, we have designed and implemented 6 configurable components to support distributed real-time aperiodic and periodic end-to-end tasks, using ACE/TAO/CIAO version 5.6/1.6/0.6. The Task Effector

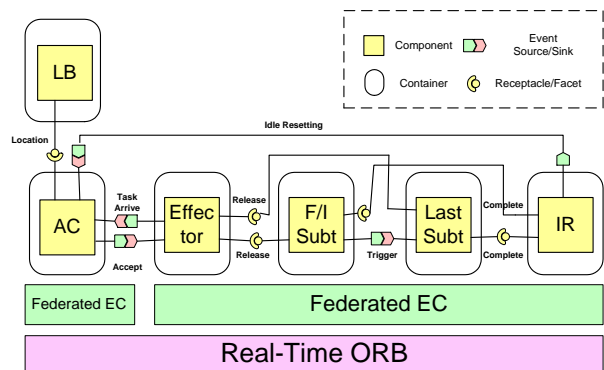


Figure 3. Component Implementation

(TE) component holds the arriving tasks, waits for the AC component decision and releases tasks. The Admission Control (AC) component decides whether to accept tasks. The Load Balancing (LB) component computes task allocations so as to balance the processors’ synthetic utilizations. The First/Intermediate (F/I) Subtask component executes the first or an intermediate subtask at a given priority. The Last Subtask component executes the last subtask at a given priority. The Idle Resetting (IR) component reports

the completed subtasks when a processor goes idle.

Each component may have several configurable attributes, so that it can be instantiated with different properties, like its criticality and execution time (for application components) or its strategy (for AC, IR and LB components). As we discussed in Section 3, our admission control and load balancing approaches adopt a centralized architecture, which employs one AC component instance and one LB component instance running on a central processor (called the “Task Manager” processor).

Each application processor contains one instance of a TE component and one instance of an IR component. The TE component on each processor reports the arrival of tasks on that processor to the AC component, which then releases or rejects the tasks based on the admission control decision. Each end-to-end task is implemented by a chain of F/I Subtask components and one Last Subtask component. We now describe the behavior of each component in detail.

Task Effector (TE) Component: When a task arrives, the TE component puts it into a waiting queue and pushes a “Task Arrive” event to the AC component. When the TE component receives an “Accept” event from the AC component, the corresponding task waiting in the queue will be released immediately. The TE component has two configurable attributes. One is a processor ID, which distinguishes TE component instances deployed on different processors. The other is the Per-job/Per-task attribute, which indicates whether before releasing any job of a periodic task the TE component will hold it until receiving an “Accept” event from the AC component. If the attribute is set to be Per-task, when a periodic task is admitted, all subsequent jobs from that periodic task can be released immediately. These attributes can be set at the creation of a TE component instance and also may be modified at run-time.

First/Intermediate (F/I) and Last Subtask Components: Both the F/I and Last Subtask components execute application subtasks. The only difference between these two kinds of components is that the F/I Subtask component has an extra port that publishes “Trigger” events to initiate the execution of the next subtask. The Last Subtask component does not need this port, since the last subtask does not have a next subtask. Each instance of these kinds of components contains a dispatching thread that executes a particular subtask at a specified priority. Both kinds of components have three configurable attributes. The first two attributes are the task execution time and priority level, which are normally set at the creation of the component instances as specified by application developers. The third attribute is No-IR, IR-per-task, or IR-per-job, which means the resetting rule either is not enabled or is enabled per task or per job respectively. Per-task means the Idle Resetting component will not be notified when periodic subjobs complete. Since each job of an aperiodic task can be treated as an independent

aperiodic task with one release, the idle resetting component is notified when aperiodic subjobs complete. Both F/I Subtask and Last Subtask components call the “Complete” method of the local IR component instance when needed. The dispatching threads in a F/I Subtask or a Last Subtask component are triggered by either a “Release” method call from the local TE component instance or a “Trigger” event from a previous F/I Subtask component instance.

Idle Resetting (IR) Component: It receives “Complete” method calls from local F/I or Last Subtask components, and pushes “Idle Resetting” events to the AC component. It has one attribute, the processor ID, which distinguishes component instances sitting on different processors.

Admission Control (AC) Component: It consumes “Task Arrive” events from the TE components and “Idle Resetting” events from the IR components. It publishes “Accept” events to the TE components to allow task releases. It makes “Location” method calls on the LB component to ask for proposed task assignment plans. The AC component has a No-LB/LB-per-task/LB-per-job attribute, which indicates whether load balancing is enabled, and if it is enabled whether it is per task or per job. If that attribute is set to LB-per-task, once a periodic task is admitted its subtask assignment is decided and kept for all following jobs. However, aperiodic tasks do not have this restriction as they are only allocated at their single job arrival time. A value of LB-per-job means the subtask assignment plan can be changed for each job of an accepted task.

Load Balancing (LB) Component: It receives “Location” method calls from the AC component, which ask for assignment plans for particular tasks. The LB component tries to balance the synthetic utilization among all processors, and may modify a previous allocation plan for a task when a new job of the task arrives. It returns an assignment plan that is acceptable and attempts to minimize the differences among synthetic utilizations on all processors after accepting that task. Alternatively, the LB component may tell the AC component that the system would be unschedulable if the task were accepted.

6 Deployment and Configuration

While our configurable components is a step towards flexible middleware services for handling aperiodic and periodic events, CPS developers still face the challenges of choosing the best combinations of strategies, assembling and deploying the components, which are tedious and error-prone if performed by hand. Therefore, we have developed a tool that automates the selection, deployment, and configuration of these components. Our tool has two key advantages: (1) it allows application developers to specify the characteristics of the CPS and automatically map

them to suitable middleware strategies, and (2) it identifies incorrect combinations of service strategies to prevent erroneous middleware configurations. CIAO’s realization of the OMG’s Light Weight Deployment and Configuration specification [21] is called the Deployment and Configuration Engine (DAnCE) [9]. DAnCE can translate an XML-based assembly specification into the execution of deployment and configuration actions needed by an application. Assembly specifications are encoded as descriptors which describe how to build cyber-physical systems using available component implementations. Information contained in the descriptors includes the number of processors, what component implementations to use, how and where to instantiate components, and how to connect component instances in an application.

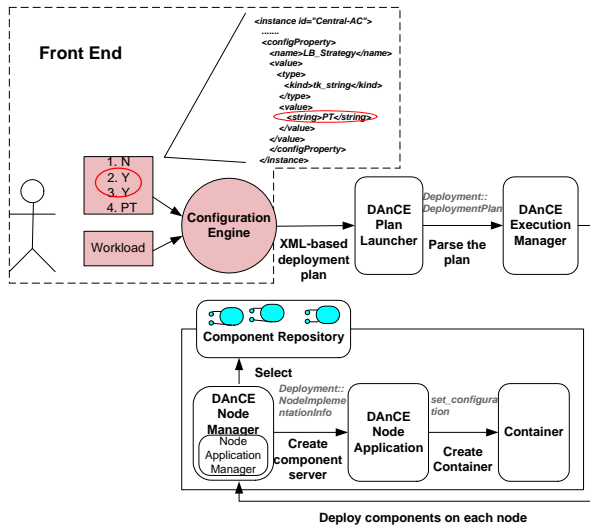


Figure 4. Dynamic Configuration Process

Front-end Configuration Engine: Although tools such as CoSMIC [12] are provided to help generate those XML files, those tools do not consider the configuration requirements of the new services we have created. We therefore provide a specific configuration engine that acts as a front-end to DAnCE, to configure our services for application developers who require configurable aperiodic scheduling support. This extension to DAnCE helps to alleviate complexities associated with deploying and configuring our services. The application developer first provides a workload specification file which describes each end-to-end task and where its subtasks execute. Then our front-end configuration engine allows the application developer to specify the characteristics of the CPS.

- (1) Does your application allow job skipping?
- (2) Does your application have replicated components?
- (3) Does your application require state persistence?
- (4) How much extra overhead can you accept as it potentially improves schedulability? [none (N), some per task

(PT), some per job (PJ)]

The front-end configuration engine parses the workload specification file and automatically maps application characteristics specified by the developer to proper configuration settings for the admission control, idle resetting and load balancing services. Finally an XML-based deployment plan is generated, which can be recognized by DAnCE. As an example, Figure 4 shows an example set of answers to those four questions. Based on the answers, the AC, IR and LB services should all be configured using per-task (PT) strategy. Figure 4 also shows part of the XML file generated by our configuration engine, with the LB strategy (LB_Strategy) setting of PT, which is due to the developer’s answers to second and third questions.

To enforce end-to-end deadline monotonic scheduling, the F/I Subtask and Last Subtask components both expose an attribute called “priority”. When our configuration engine reads the workload specification file, it assigns priorities in order of tasks’ end-to-end deadlines, and writes this priority information into the generated XML deployment plan, to be parsed by DAnCE later. Our front-end configuration engine not only generates well formed assembly specifications, according to the application developers’ instructions, but it also performs a feasibility check on configuration settings, to ensure correct handling of dependent constraints. For example, per task admission control with per job idle resetting would be contradictory as we mentioned in Section 4. Since a developer might specify incompatible service configuration combinations, our approach should be able to detect and disallow them. If application characteristics are not provided by the developers, our configuration engine can supply default configuration settings, i.e., per task admission control, idle resetting and load balancing.

We have used the <configproperty> feature of DAnCE to extend the set of attributes that can be configured flexibly according to other configuration settings. For example, if the load balancing service is configured using the per-task strategy, the corresponding property of the AC component should also be set to per-task. DAnCE’s Plan Launcher parses the XML-based deployment plan and stores the property name (LB_Strategy) and value in a data structure (Property) which is a field of the AC instance definition structure. The definitions of the AC instance and all other component instances comprise a deployment plan (Deployment::DeploymentPlan) that is then passed to DAnCE’s Execution Manager for execution. The Execution Manager propagates the deployment plan data structure to DAnCE’s Node Application Manager, which parses it into an initialization data structure (NodeImplementationInfo). Finally, the Node Application Manager passes the initialization data structure to the Node Application. When the Node Application installs the AC component instance, it also initializes the LB_Strategy attribute of the AC component through

a standard Configurator interface (`set_configuration`), using the initialization data structure it received.

7 Experimental Evaluations

To validate our approach, and to evaluate the performance, overheads and benefits resulting from it, we conducted a series of experiments which we describe in this section. The experiments were performed on a testbed consisting of six machines connected by a 100Mbps Ethernet switch. Two are Pentium-IV 2.5GHz machines with 1G RAM and 512K cache each, two are Pentium-IV 2.8GHz machines with 1G RAM and 512K cache each, and the other two are Pentium-IV 3.4GHz machines with 2G RAM and 2048K cache each. Each machine runs version 2.4.22 of the KURT-Linux operating system. One Pentium-IV 2.5GHz machine is used as a central task manager where the AC and LB components are deployed. The other five machines are used as application processors on which TE, F/I Subtask, Last Subtask and IR components are deployed.

7.1 Random Workloads

We first randomly generated 10 sets of 9 tasks, each including 4 aperiodic tasks and 5 periodic tasks. The number of subtasks per task is uniformly distributed between 1 and 5. Subtasks are randomly assigned to 5 application processors. Task deadlines are randomly chosen between 250 ms and 10 s. The periods of periodic tasks are equal to their deadlines. The arrival of aperiodic tasks follows a Poisson distribution. The synthetic utilization of every processor is 0.5, if all tasks arrive simultaneously. Each subtask is assigned to a processor, and has a duplicate sitting on a different processor which is randomly picked from the other 4 application processors.

In this experiment, we evaluated all 15 reasonable combinations of strategies, since it is convenient to choose and run different combinations with the help of our configuration engine. We ran 10 task sets using each combination and compared them. The performance metric we used in these evaluations is the *accepted utilization ratio*, i.e., the total utilization of jobs actually released divided by the total utilization of all jobs arriving. To be concise, we use capital letters to represent strategies: **N** when a service is **not** enabled in this configuration; **T** when a service is enabled for each **task**; and **J** when a service is enabled for each **job** of a task. In the following figures, a three element tuple denotes each combination of settings for the three configurable services: first for the admission control service, then for the idle resetting service, and last for the load balancing service.

The bars in Figure 5 show the average (mean) results over the 10 task sets. As is shown in Figure 5, enabling




Figure 5. Accepted Utilization Ratio

either idle resetting or load balancing can increase the utilization of tasks admitted. Moreover, the experiment shows that enabling IR per job (`*_J_*`) significantly outperforms the configurations which enable IR per task (`*_T_*`) or not at all (`*_N_*`). This is because IR per job removes the contribution of all completed periodic subjobs to the synthetic utilizations which greatly helps to admit more jobs. Enabling all three services per job (`J_J_J`) performed comparably to the other (`J_J_*`) configurations (averaging higher though the differences were not significant) and outperformed all other configurations significantly, even though the `J_J_J` configuration introduces the most overhead. We also notice the difference is small when we only change the configuration of the LB component and keep the configuration of other two services the same. This is because when we randomly generated these 10 task sets, the resulting synthetic utilization of each processor was similar. To show the potential benefit of the LB component, we designed another experiment that is described in the next section.

7.2 Imbalanced Workloads

In the second experiment, we use an imbalanced workload. It is representative of a dynamic CPS in which a subset of the system processors may experience heavy load. For example, in an industrial control system, a blockage in a fluid flow valve may cause a sharp increase in the load on the processors immediately connected to it, as aperiodic alert and diagnostic tasks are launched. In this experiment, we divided the 5 application processors into two groups. One group contains 3 processors hosting all tasks. The other group contains 2 processors hosting all duplicates. 10 task sets are randomly generated as in the above experiment, except that all subtasks were randomly assigned to 3 application processors in the first group and the number of subtasks per task is uniformly distributed between 1 and 3. The synthetic utilization for any of these three processors is 0.7.

Each subtask has one replica sitting on one processor in the second group.

Each of 10 task sets was run for the 15 different valid combinations, and for each combination we then averaged the utilization acceptance ratio over the 10 results. These 15 combinations can be divided into 5 sets. Each set contains three combinations represented by three adjacent bars in Figure 6. In each set, we kept the admission control and idle resetting strategies the same, but changed the load balancing strategy from none to per task, then to per job. As figure 6 shows, load balancing per task provides a significant improvement when compared with the results without load balancing. However, there is not much difference between load balancing per task vs. per job.




Figure 6. LB Strategy Comparison

From these two experiments, we found that configuring different strategies according to application characteristics can have a significant impact on the performance of a cyber-physical system with aperiodic and periodic events. Our design of the AC, IR and LB services as easily configurable components allows application developers to explore and select valid configurations based on the characteristics and requirements of their applications, and based on the trade-offs indicated by these empirical results.

7.3 Overheads of Service Components

To evaluate the efficiency of our component-based middleware services, we measured overheads using 3 of the processors to run application components and another processor to run the AC and LB components. The workload is randomly generated in the same way as described in Section 7.1, except that the number of subtasks per task is uniformly distributed between 1 and 3. Each experiment ran for 5 minutes. We examined the different sources of overhead that may occur when a task arrives at TE component TE1, after which AC and LB components run the task in component TE1 or re-allocate it to another TE component, TE2. Figure 7 shows how the total delay for each service

includes the costs of operations located in several components. Figure 8 lists the operation numbers shown in Figure 7 to provide a detailed accounting of the delays resulting from different combinations of service configurations.

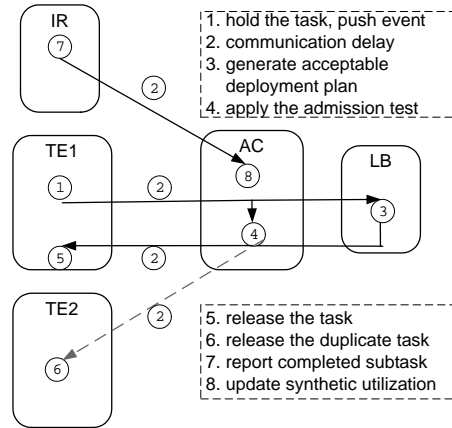


Figure 7. Sources of Overhead/Delay

	mean	max
AC without LB (1+2+4+2+5)	1114	1248
AC with LB (1+2+3+2+5) (no re-allocation)	1116	1253
AC with LB (1+2+3+2+6) (re-allocation)	1201	1327
LB (no re-allocation) (1+2+3+2+5)	1113	1250
LB (re-allocation) (1+2+3+2+6)	1198	1319
IR (on AC side) (8)	17	18
IR (other part) (7+2)	662	683
Communication Delay (2)	322	361

Figure 8. Service Overheads (μ s)

To calculate the delays for AC without LB, AC with LB without re-allocation and LB without re-allocation, we can simply calculate the interval between when one task arrives on a processor and when the task is released on the processor. However, if the LB component re-allocates the first subtask on a different processor using its duplicate, as in the case of AC with LB, it is difficult to determine a precise time interval between when one task arrives on one processor and when it is released on another processor, because our experiment environment does not provide sufficiently high resolution time synchronization among processors, which is an inherent limitation for many distributed cyber-physical systems. We therefore measure the overheads on all involved processors individually, then add them together plus twice the communication delays (step 2 in Figure 7) between the processors. Three processors are involved: the processor

where the task arrives (step 1), the central task manager processor (steps 3 and 4) and the processor where the duplicate task is released (step 6). We ran this experiment using KURT-Linux version 2.4.22, which provides a CPU-supported timestamp counter with nanosecond resolution. By using instrumentation provided with the KURT-Linux distribution, we can obtain a precise accounting of operation start and stop times and communication delays. To measure the communication delay between the application processor and the admission control processor on our experimental platform, we pushed an event back and forth between the application processor to the admission control processor 1000 times, then calculated the mean and max value among 1000 results. We then divided the round trip time by 2 to obtain the approximate mean and maximum communication delays between the application processor and the admission control processor.

The total delay for LB when reallocation happens, is measured in the same way as for the case of AC with LB with reallocation. To calculate the delay from the IR component, we divide its execution into two parts. The small overhead on the admission control processor must be counted in the overall delay. The large overhead on the application processor and the communication delay only happen during CPU idle time, and although it represents an additional overhead induced by the IR component, it does not affect performance, which is why we report the two parts separately in Figure 8. From the results in Figure 8, we can see that all of the delays induced by our configurable components are less than 2 ms, which is acceptable for many distributed cyber-physical systems. For applications with tight schedules, a developer can make further decisions on how to configure services based on this delay information and based on the effects of the different configurations on task management, which we discussed in Section 4.

8 Related Work

Component Middleware: The architectural patterns used in the CORBA Component Model (CCM) [26] are also used in other popular component middleware technologies, such as J2EE [27, 3]. Among the existing component middleware technologies, CCM is the most suitable for distributed cyber-physical systems since CORBA is the only standards-based COTS middleware that explicitly considers the real-time requirements of distributed cyber-physical systems.

QoS-aware middleware: Quality Objects (QuO) [23, 31] is an adaptive middleware framework developed by BBN Technologies that allows developers to use aspect-oriented software development techniques to separate the concerns of QoS programming from application logic in distributed cyber-physical systems. A Qosket is a unit of encapsulation and reuse for QuO systemic behaviors. In comparison

to CIAO, Qoskets and QuO emphasize dynamic QoS provisioning where CIAO emphasizes static QoS provisioning and integration of various mechanisms and behaviors during different stages of the development lifecycle. The dynamic-TAO [16] project applies reflective techniques to reconfigure Object Request Broker (ORB) components at run-time. Similar to dynamicTAO, the Open ORB [11] project also aims at highly configurable and dynamically reconfigurable middleware platforms to support applications with dynamic requirements. Zhang et al. [29] use aspect-oriented techniques to improve the customizability of the middleware core infrastructure at the ORB level.

QoS-aware component Middleware: Component middleware's container architecture enables meta-programming of QoS attributes in component middleware. For example, aspect-oriented techniques can be used to plug in different systemic behaviors [6]. This approach is similar to CIAO in that it provides mechanisms to inject aspects into systems at the middleware level. de Miguel's work [8] further develops the state of the art in QoS-enabled containers by extending a QoS EJB container interface to support a QoSContext interface that allows the exchange of QoS-related information among component instances. To take advantage of the QoS-container, a component must implement QoSBean and QoSNegotiation interfaces. However, this requirement increases dependence among component implementations. The QoS Enabled Distributed Objects (Qedo) [10] project is another effort to make QoS support an integral part of CCM. Qedo's extensions to the CCM container interface and Component Implementation Framework (CIF) require component implementations to interact with the container QoS interface and negotiate the level of QoS contract directly. Although this approach is suitable for certain applications where QoS is part of the functional requirements, it tightly couples the QoS provisioning and adaptation behaviors into the component implementation, which may limit the reusability of the component. In comparison, CIAO explicitly avoids this coupling and composes the QoS aspects into applications declaratively. There have been several other efforts to introduce of QoS in conventional component middleware platforms. The FIRST Scheduling Framework (FSF) [2] proposes to compose several applications and to schedule the available resources flexibly while guaranteeing hard real-time requirements. A real-time component type model [7], which integrates QoS facilities into component containers also was introduced, based on the EJB and RMI specifications. A schedulability analysis algorithm [18] for hierarchical scheduling systems has been introduced for dependent components which interact through remote procedure calls. None of these approaches provides the configurable services for mixed aperiodic and periodic end-to-end tasks offered by our approach.

9 Conclusions

The work presented in this paper represents a promising step towards configurable admission control and load balancing support for a variety of distributed cyber-physical systems with aperiodic and periodic events. We first analyzed a set of key characteristics of different cyber-physical systems and mapped them to suitable strategies for middleware services. We then have designed and implemented configurable middleware components that provide effective on-line admission control and load balancing and can be easily configured and deployed on different processors. Our front-end configuration engine can automatically process the user's specification file and generate a corresponding deployment plan for DAnCE, thus making it easier for developers to select suitable configurations, and to avoid invalid ones. Empirical results we obtained showed that (1) our configurable component middleware is well suited for satisfying different applications with a variety of alternative characteristics and requirements, and (2) our component middleware services are appropriately efficient on a Linux platform.

References

- [1] T. F. Abdelzaher, G. Thaker, and P. Lardieri. A Feasible Region for Meeting Aperiodic End-to-end Deadlines in Resource Pipelines. In *ICDCS*, 2004.
- [2] M. Aldea, G. Bernat, I. Broster, A. Burns, R. Dobrin, J. M. Drake, G. Fohler, P. Gai, M. G. Harbour, G. Guidi, J. J. Gutiérrez, T. Lennvall, G. Lipari, J. M. Martínez, J. L. Medina, J. C. Palencia, and M. Trimarchi. FSF: A Real-Time Scheduling Architecture Framework. In *RTAS*, 2006.
- [3] D. Alur, J. Crupi, and D. Malks. *Core J2EE Patterns: Best Practices and Design Strategies*. Prentice Hall, 2001.
- [4] B. Andersson and C. Ekelin. Exact Admission-Control for Integrated Aperiodic and Periodic Tasks. In *RTAS*, 2005.
- [5] G. Bollella, J. Gosling, B. Brosgol, P. Dibble, S. Furr, D. Hardin, and M. Turnbull. *The Real-Time Specification for Java*. Addison-Wesley, 2000.
- [6] D. Conan, E. Putrycz, N. Farcet, and M. DeMiguel. Integration of Non-Functional Properties in Containers. *Proceedings of the Sixth International Workshop on Component-Oriented Programming (WCOP)*, 2001.
- [7] M. A. de Miguel. Integration of QoS Facilities into Component Container Architectures. In *ISORC*, 2002.
- [8] M. A. de Miguel. QoS-Aware Component Frameworks. In *The 10th International Workshop on Quality of Service (IWQoS 2002)*, Miami Beach, Florida, May 2002.
- [9] G. Deng, D. C. Schmidt, C. Gill, and N. Wang, editors. *QoS-Enabled Component Middleware for Distributed Real-Time and Embedded Systems*. CRC Press, 2007.
- [10] FOKUS. Qedo Project Homepage. <http://qedo.berlios.de/>.
- [11] G. S. Blair and G. Coulson and A. Andersen and L. Blair and M. Clarke and F. Costa and H. Duran-Limon and T. Fitzpatrick and L. Johnston and R. Moreira and N. Parlavantzis and K. Saikoski. The Design and Implementation of Open ORB 2. *IEEE Distributed Systems Online*, 2(6), June 2001.
- [12] A. Gokhale. Component Synthesis using Model Integrated Computing. www.dre.vanderbilt.edu/cosmic, 2003.
- [13] T. H. Harrison, D. L. Levine, and D. C. Schmidt. The design and performance of a real-time CORBA event service. In *Proceedings of OOPSLA*, 1997.
- [14] Institute for Software Integrated Systems. Component-Integrated ACE ORB (CIAO). www.dre.vanderbilt.edu/CIAO/, Vanderbilt University.
- [15] Institute for Software Integrated Systems. The ACE ORB (TAO). www.dre.vanderbilt.edu/TAO/, Vanderbilt University.
- [16] F. Kon, F. Costa, G. Blair, and R. H. Campbell. The Case for Reflective Middleware. *Commun. ACM*, 45(6):33–38, June 2002.
- [17] G. Koren and D. Shasha. Skip-over: Algorithms and complexity for overloaded systems that allow skips. In *RTSS*, 1995.
- [18] J. L. Lorente, G. Lipari, and E. Bini. A Hierarchical Scheduling Model for Component-Based Real-Time Systems. In *WPDRTS*, 2006.
- [19] Object Management Group. *CORBA Components*, OMG Document formal/2002-06-65 edition, June 2002.
- [20] Object Management Group. *Real-Time CORBA Specification*, 1.1 edition, Aug. 2002.
- [21] Object Management Group. *Light Weight CORBA Component Model Revised Submission*, OMG Document realtime/03-05-05 edition, May 2003.
- [22] F. H. Raven. *Automatic Control Engineering*. Mcgraw-Hill, New York, New York, fifth edition, 1994.
- [23] R. Schantz, J. Loyall, M. Atighetchi, and P. Pal. Packaging Quality of Service Control Behaviors for Reuse. In *ISORC*, pages 375–385, April/May 2002.
- [24] L. Sha and et. al. Real Time Scheduling Theory: A Historical Perspective. *The Journal of Real-Time Systems*, 10:101–155, 2004.
- [25] J. Strosnider, J. P. Lehoczky, and L. Sha. The deferrable server algorithm for enhanced aperiodic responsiveness in real-time environments. *IEEE Transactions on Computers*, 44(1):73–91, 1995.
- [26] M. Volter, A. Schmid, and E. Wolff. *Server Component Patterns – Component Infrastructures illustrated with EJB*. Wiley & Sons, New York, 2002.
- [27] M. Volter, A. Schmid, and E. Wolff. *Server Component Patterns: Component Infrastructures Illustrated with EJB*. Wiley Series in Software Design Patterns, West Sussex, England, 2002.
- [28] N. Wang, C. Gill, D. C. Schmidt, and V. Subramonian. Configuring Real-time Aspects in Component Middleware. In *Proc. of the International Symposium on Distributed Objects and Applications (DOA'04)*, Agia Napa, Cyprus, Oct. 2004.
- [29] C. Zhang and H.-A. Jacobsen. Resolving Feature Convolution in Middleware Systems. In *OOPSLA*, 2004.

- [30] Y. Zhang, C. Lu, C. Gill, P. Lardieri, and G. Thaker. Middle-ware Support for Aperiodic Tasks in Distributed Real-Time Systems. In *RTAS*, 2007.
- [31] 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.