

Provided for non-commercial research and educational use only.  
Not for reproduction or distribution or commercial use.



This article was originally published in a journal published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues that you know, and providing a copy to your institution's administrator.

All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<http://www.elsevier.com/locate/permissionusematerial>



ELSEVIER

Available online at [www.sciencedirect.com](http://www.sciencedirect.com)

 ScienceDirect

The Journal of Systems and Software 80 (2007) 938–950

 **The Journal of  
Systems and  
Software**

[www.elsevier.com/locate/jss](http://www.elsevier.com/locate/jss)

# FC-ORB: A robust distributed real-time embedded middleware with end-to-end utilization control <sup>☆</sup>

Xiaorui Wang <sup>a,\*</sup>, Yingming Chen <sup>b</sup>, Chenyang Lu <sup>b</sup>, Xenofon Koutsoukos <sup>c</sup>

<sup>a</sup> *University of Tennessee, Knoxville, United States*

<sup>b</sup> *Washington University in St. Louis, United States*

<sup>c</sup> *Vanderbilt University, United States*

Available online 1 November 2006

## Abstract

A key challenge for distributed real-time and embedded (DRE) middleware is maintaining both system reliability and desired real-time performance in unpredictable environments where system workload and resources may fluctuate significantly. This paper presents FC-ORB, a real-time Object Request Broker (ORB) middleware that employs end-to-end utilization control to handle fluctuations in application workload and system resources. The contributions of this paper are three-fold. First, we present a novel utilization control service that enforces desired CPU utilization bounds on multiple processors by adapting the rates of end-to-end tasks within user-specified ranges. Second, we describe a set of middleware-level mechanisms designed to support end-to-end tasks and distributed multi-processor utilization control in a real-time ORB. Finally, we present extensive experimental results on a Linux testbed. Our results demonstrate that our middleware can maintain desired utilizations in face of uncertainties and variations in task execution times, resource contentions from external workloads, and permanent processor failure. FC-ORB demonstrates that the integration of utilization control, end-to-end scheduling, and fault-tolerance mechanisms in DRE middleware is a promising approach for enhancing the robustness of DRE applications in unpredictable environments.

© 2006 Elsevier Inc. All rights reserved.

*Keywords:* Real-time systems; Middleware; Embedded systems; Distributed systems; Feedback control

## 1. Introduction

Distributed real-time and embedded (DRE) applications have stringent requirements for end-to-end timeliness and reliability whose assurance is essential to their proper operation. In recent years, many DRE systems have become open to unpredictable operating environments where both system workload and platform may vary significantly at run time. For example, the execution of data-driven applications such as autonomous surveillance is heavily influenced by sensor readings. External events such as

detection of an intruder can trigger sudden increase in system workloads. Furthermore, many mission-critical applications must continue to provide real-time services despite hardware failures, software faults, and cyber attacks.

While DRE middleware has shown promise in improving the real-time properties of many applications, existing middleware systems often do not work well in unpredictable environments due to their dependence on traditional real-time schedulability analysis. When accurate knowledge about workloads and platforms is not available, a DRE application configured based on schedulability analysis may suffer deadline misses or even system crash (Lu et al., 2003). A critical challenge faced by application developers is to achieve *robust* real-time performance in unpredictable environments. Since in DRE systems, an end-to-end application that violates its real-time properties is equivalent to (or sometimes even worse than) an

<sup>☆</sup> Parts of this work were presented at a conference paper Wang et al., 2005. This research was supported in part by NSF CAREER award (Grant CNS-0448554) and DARPA Adaptive and Reflective Middleware Systems (ARMS) program (Grant NBCHC030140).

\* Corresponding author. Tel.: +1 865 974 0627; fax: +1 865 974 5483.  
E-mail address: [xwang@ece.utk.edu](mailto:xwang@ece.utk.edu) (X. Wang).

application that does not perform its computation, utilization guarantees affect directly the availability of the end-to-end application.

This paper presents the design and empirical evaluation of an adaptive middleware called FC-ORB (Feedback Controlled ORB) that aims to enhance the robustness of DRE applications. The novelty of FC-ORB is the integration of end-to-end scheduling, adaptive QoS control, and fault-tolerance mechanisms that are optimized for unpredictable environments. Specifically, this paper makes three contributions.

- *End-to-End Real-Time ORB*: Our ORB service supports end-to-end real-time tasks based on the end-to-end scheduling framework (Liu, 2000). The FC-ORB architecture is designed to facilitate efficient end-to-end adaptation and fault-tolerance in memory-constrained DRE systems.
- *End-to-End Utilization Control*: The utilization control service enforces desired CPU utilizations in a DRE system despite significant uncertainties in system workloads. The core of the utilization control service is a distributed feedback control loop that coordinates adaptations on multiple interdependent processors.
- *Adaptive Fault Tolerance*: FC-ORB handles processor failures with an adaptive strategy that combines reconfigurable utilization control and task migration. A unique feature of our fault tolerance approach is that it can maintain *real-time* properties for DRE applications after a processor failure.

FC-ORB has been implemented and evaluated on a Linux platform. Our experimental results demonstrate that FC-ORB can significantly improve the end-to-end real-time performance of DRE middleware in face of a broad set of dynamic uncertainties and fluctuations in task execution times, resource contention from external workloads, and processor failures. FC-ORB demonstrates that the integration of utilization control, end-to-end scheduling, and fault-tolerance mechanisms in DRE middleware is a promising approach for enhancing the robustness of DRE applications in unpredictable environments.

FC-ORB is particularly useful for DRE applications that are amenable to rate adaptation such as digital feedback control systems (Marti et al., 2002; Seto et al., 1996), monitoring systems (Zhao et al., 2001), and multimedia (Brandt et al., 1998). In these systems, task rates can be adjusted without causing system failure. Furthermore, tasks running at higher rates contribute higher values to the application (e.g., increasing the sampling rate of a digital controller improves the control performance). Our framework can benefit Supervisory Control and Data Acquisition (SCADA) systems which provide monitoring and control functions that are inherently periodic at geographically distributed sites.

The rest of the paper is organized as follows. Section 2 describes the design of the FC-ORB architecture. Section

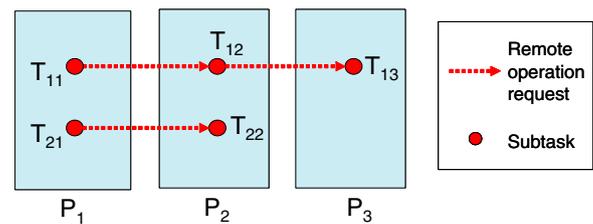


Fig. 1. An example DRE application.

3 presents the experimental results. Section 4 highlights the contributions of FC-ORB by comparing it with related works. Section 5 concludes the paper.

## 2. Design of the FC-ORB architecture

In this section, we first introduce the end-to-end task model and scheduling framework supported by FC-ORB. We then describe the main components of FC-ORB: the end-to-end ORB service, the utilization control service, and the adaptive fault-tolerance mechanisms.

### 2.1. Applications

FC-ORB supports an end-to-end task model (Liu, 2000) employed by many DRE applications. An application is comprised of  $m$  periodic tasks  $\{T_i | 1 \leq i \leq m\}$  executing on  $n$  processors  $\{P_i | 1 \leq i \leq n\}$ . Task  $T_i$  is composed of a chain of subtasks  $\{T_{ij} | 1 \leq j \leq n_i\}$  which are implemented as a sequence of object operations on different processors.<sup>1</sup> A subtask may be executed by one or more operation requests on a same processor. The invocation of a subtask  $T_{ij} (1 < j \leq n_i)$  is triggered by its predecessor  $T_{i,j-1}$  through a remote operation request. A non-greedy synchronization protocol called release guard (Sun and Liu, 1996) is used to ensure that the interval between two consecutive releases of the same subtask is not less than its period. Hence, all the subtasks of a periodic task share the same rate as the first subtask. In FC-ORB, the rate of a task (and all its subtasks) can be adjusted by changing the rate of its first subtask. An example DRE application with two end-to-end tasks running on three processors is shown in Fig. 1.

Our application model has two important properties. First, while each subtask  $T_{ij}$  has an *estimated* execution time  $c_{ij}$  available at design time, its *actual* execution time may be different from its estimation and may vary at run-time. Such uncertainty is common for DRE systems operating in unpredictable environments. Second, the rate of a task  $T_i$  may be dynamically adjusted within a range  $[R_{\min,i}, R_{\max,i}]$ . This assumption is based on the fact that the task rates in many DRE applications (e.g., digital control (Marti et al., 2002; Seto et al., 1996), sensor update, and multimedia (Brandt et al., 1998)) can be dynamically adjusted without causing system failure. A task running

<sup>1</sup> FC-ORB can be extended to support a more general task model in which a task is composed of a graph of subtasks (Liu, 2000).

at a higher rate contributes a higher value to the application at the cost of higher utilization. For instance, although a digital control system usually has better control performance when it executes at a higher rate, it can usually remain stable when executing at a lower rate.

Each task  $T_i$  is subject to an end-to-end soft deadline related to its period. FC-ORB implements the end-to-end scheduling approach (Sun and Liu, 1996) to meet task deadlines. The deadline of a task is divided into subdeadlines of its subtasks (Kao and Garcia-Molina, 1997; Natale M.D. and Stankovic, 1994). The release guard protocol is used to synchronize the execution of subtasks such that each subtask can be modeled as a periodic task. Hence, the problem of meeting the deadline is transformed to the problem of meeting the subdeadline of each subtask. A well known approach for meeting the subdeadlines on a processor is to ensure that its utilization remains below its schedulable utilization bound (Lehoczky, 1990; Liu and Layland, 1973). Therefore the end-to-end scheduling approach enables FC-ORB to meet end-to-end deadlines by controlling the utilizations of all processors in the system.

## 2.2. Middleware support for end-to-end tasks

### 2.2.1. Implementation of end-to-end tasks

Fig. 2 illustrates the FC-ORB implementation of the example DRE application shown in Fig. 1. Each subtask is executed by a separate thread whose priority is decided by a priority manager. In Fig. 2, each dashed box spanning from the application layer to the ORB core layer represents a subtask in Fig. 1. Every subtask is associated with a separate Reactor (Schmidt, 1995) to create timeout events and to manage communication connections.

As shown in Fig. 2, the first subtask of a task is implemented with a periodic ACE timer, a reactor, and a connector (Schmidt, 1997). The ACE Connector framework is used to decouple communication initialization from application-specific tasks that communication services perform once initialization is complete. Connector can be configured with different IPC mechanisms to support communication in different distributed applications. The ACE Reactor framework is an extensible, object-oriented demultiplexer

that dispatches events to application-specific handlers. It can support I/O-based, timer-based, signal-based, and synchronization-based events. It simplifies the development of event-driven programs for many distributed applications. The timer periodically triggers a local operation (e.g., a method of an object) which implements the functionality of this subtask. Following the execution of this operation, a one-way remote operation request is pushed through the Connector to the succeeding subtask that is located on another processor. The succeeding subtask employs an Acceptor (Schmidt, 1997) to accept the request from its preceding subtask. Each pair of connector and acceptor maintains a separate TCP connection to avoid priority inversion in the communication subsystem. The release guard protocol enforces to be the interval between two successive invocations of a same subtask is bounded below by its period. Earlier research has shown that the release guard protocol can effectively reduce the end-to-end response time and jitter of tasks in DRE systems (Sun and Liu, 1996). FC-ORB implements the release guard protocol with a FIFO waiting queue and one-shot ACE timers. Upon receiving a remote operation request, a subtask compares the current time with the last invocation time of this operation. Based on the release guard rules (Sun and Liu, 1996), the subtask either immediately invokes the requested operation or enqueues this request to the waiting queue if the request arrives too early. When the request is enqueued, a one-shot ACE timer is registered with the reactor to trigger the requested operation at the time that equals the last invocation time plus the task's period. After the one-shot timer fires and the enqueued request is served, a remote operation request is sent to the next subtask in the end-to-end task chain. An end-to-end real-time task is completed when the execution of its last subtask is completed.

### 2.2.2. Priority management

The integration of end-to-end scheduling and utilization control introduces new challenges to the design of scheduling mechanisms in ORB middleware. For instance, the rate adaptation mechanism adopted by FC-ORB and several other projects (Lu et al., 2003; Lu et al., 2005) may dynamically change the rates of end-to-end tasks. This may cause

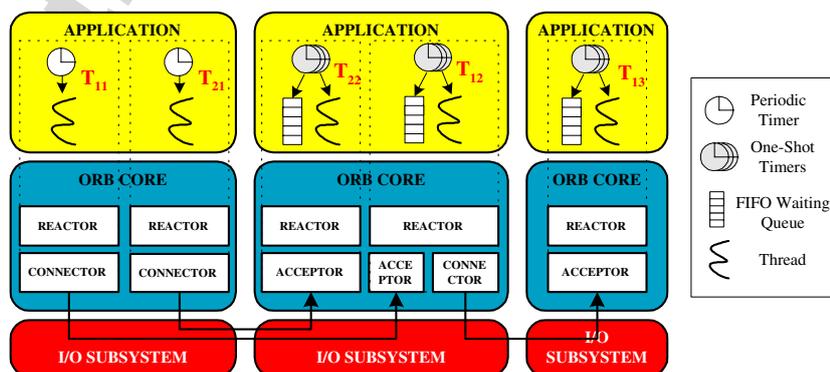


Fig. 2. FC-ORB's end-to-end ORB architecture.

the middleware to change the priorities of all its subtasks, e.g., when the rate monotonic scheduling (RMS) policy is used. To satisfy the special requirements posed by rate adaptation and end-to-end scheduling, our ORB service is configured with the *server-declared priority* model (Schmidt and Kuhns, 2000) and the *thread-per-subtask* concurrency architecture.

To support the server-declared priority model, FC-ORB implements a priority manager on each processor to assign priorities to local subtasks. The incoming requests from another processor are served by a thread with a real-time priority dictated by the priority manager located on the host processor. Currently the priority manager only supports the RMS policy, although the following discussions are also applicable to other rate- or deadline-dependent scheduling policies (note that task deadlines are usually related to their periods). There are several advantages of using server-declared priority model in the FC-ORB system. First, each processor is able to change thread priorities locally, based on the current rates of the subtasks located on it, so a processor only needs to know the local subtasks. This makes the system more scalable to large applications. Moreover, the server-declared model has less overhead because it does not have to adjust a thread’s priority every time the priority of its predecessor subtask is changed, as it would do with the client-propagated model.

The thread-per-priority concurrency architecture has been adopted in existing DRE middleware (e.g., (Schmidt, 2002)). In this model, the same thread is responsible for executing all subtasks with a same priority. This is because the workload is assumed to use only a limited number of fixed task rates. However, this concurrency architecture is not suitable for rate adaptation. Due to rate adaptation, the rates and thus the priorities of subtasks vary dynamically at run-time. In such situations, the thread-per-priority architecture would require the ORB to dynamically move a subtask from one thread to another thread which can introduce significant overhead.

To avoid this problem FC-ORB implements the thread-per-subtask architecture that executes each subtask with a

separate thread. FC-ORB adjusts the priorities of the threads only when the *order* of the task rates is changed. While the task rates may vary at every control period, the order of task rates often changes at a much lower frequency. Therefore, the thread-per-subtask architecture enables FC-ORB to adapt task rates in a more flexible way, with less overhead.

A potential advantage of the thread-per-priority architecture is that it may need fewer threads to execute applications. However, as FC-ORB is targeted at memory-constrained networked embedded systems that commonly have limited number of subtasks on a processor, each subtask can be easily mapped to a thread with a unique native thread priority even in a thread-per-subtask architecture.

### 2.3. End-to-end utilization control service

FC-ORB allows users to specify a set of application parameters in a configuration file that is used to initialize the middleware when the system is started. Configuration parameters include the desired CPU utilization on each processor and the allowed range of rate for each real-time task. The utilization control service dynamically enforces the desired CPU utilizations on all processors by adapting the rates of real-time tasks within the specified ranges, despite significant uncertainties and fluctuation in system workload and platform. Therefore, to meet end-to-end deadlines, the application users only need to specify the utilization reference of each processor to a value below its schedulable utilization bound.

In the rest of this subsection, we first give an overview of the feedback control loop of the utilization control service, and then describe each component of the loop in detail.

#### 2.3.1. Feedback control loop

The utilization control service implements the EUCON algorithm (Lu et al., 2005) as a distributed feedback control loop in the middleware. As shown in Fig. 3, the feedback control loop is composed of a utilization monitor, a rate

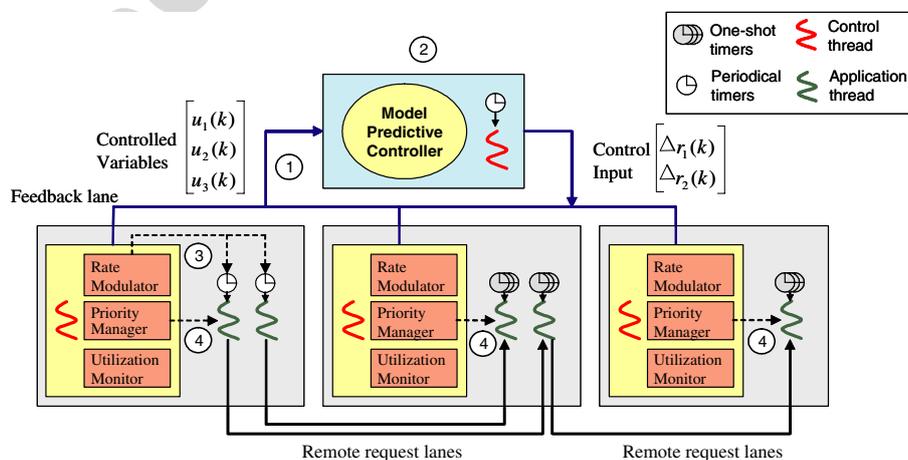


Fig. 3. The distributed feedback control loop of the utilization control service.

modulator and a priority manager on each processor, and a centralized controller.

The feedback control loop is invoked at the end of every sampling period. It works as follows: (1) the utilization monitor on each processor sends its utilization in the last sampling period to the controller; (2) the controller collects the utilizations from all processors, computes the new task rates, and sends the new task rates to the rate modulators on all processors where the tasks are running; (3) the rate modulators on processors that host the first subtasks of tasks change the rates of the first subtasks according to the input from the controller; and (4) the priority manager on each processor check and adjust the thread priorities based on the new task rates if necessary.

The controller computes the new task rates using a model predictive control (MPC) algorithm. The control algorithm solves at every time step an optimization problem that minimizes the difference between the desired and the actual utilizations subject to the task rate constraints. The optimization problem is a constraint least-square problem that can be solved efficiently using quadratic programming. Assuming that the optimization problem is feasible, i.e., there exist task rates that satisfy the utilization bounds, the stability of the controller can be formally analyzed and provide statistical guarantees for the schedulability of the system. Details can be found in Lu et al. (2005).<sup>2</sup>

As shown in Fig. 3, the three components of the feedback control loop on an application processor (i.e., a processor executing applications and the ORB) are executed by a separate thread called the *control thread*. This control thread has the highest priority in the middleware system so that the feedback control loop can be executed in overload conditions, when it is needed most. The controller is implemented as an independent process that can be deployed on a separate processor or on an application processor. The controller also serves as a coordinator of the FC-ORB system. Every application processor in the system tries to connect with the controller through a TCP connection (called *feedback lane*) when the node is started. Once all application processors are connected to the controller, the whole system starts to run the configured application.

### 2.3.2. Control components

We now present the details of each utilization control component.

**2.3.2.1. Controller.** The controller is implemented as a single-thread process. It employs a reactor to interact with all processors in the system. Each time its periodic timer fires, it sends utilization requests to all application processors through the feedback lanes. The incoming replies are registered with the Reactor as events to be handled asynchro-

nously. This enables the controller to avoid being blocked by an overloaded application processor. After it collects the replies from all processors, it executes a MPC algorithm proposed in Lu et al. (2005) to calculate the new task rates. Then, for each task whose rate needs to be changed, the controller sends the task's new rate to all processors that host one or more subtasks of the tasks whose rates have been changed. If a processor does not reply in an entire control period, its utilization is treated as 100%, as the controller assumes this processor is saturated by its workload.

**2.3.2.2. Utilization monitor.** The utilization monitor uses the `/proc/stat` file in Linux to estimate the CPU utilization in each sampling period. The `/proc/stat` file records the number of jiffies (usually 10 ms in Linux) when the CPU is in user mode, user mode with low priority (`nice`), system mode, and when used by the idle task, since the system starts. At the end of each sampling period, the utilization monitor reads the counters, and estimates the CPU utilization as 1 minus the number of jiffies used by the idle task in the last sampling period divided by the total number of jiffies in the same period.

**2.3.2.3. Rate modulator.** A Rate Modulator is located on each processor. It receives the new rates for its remote invocation requests from the controller through the feedback lane, and resets the timer interval of the first subtask of each task whose invocation rate has been changed.

**2.3.2.4. Priority manager.** All processors in FC-ORB assign priorities to their subtasks based on a real-time scheduling algorithm (e.g., RMS). It is important to strictly enforce the scheduling algorithm to achieve desired real-time performance. However, as a result of rate adaptation, a task with a rate higher than another task could be assigned a lower rate in the next sampling period. Consequently, the priority of this task has to be adjusted at run-time. The priority manager on each processor checks the rate order of all subtasks on this processor. If the rate order of two or more subtasks is reversed, the priority manager reassigns the correct priorities for the threads of those subtasks.

## 2.4. Fault tolerance

A robust DRE middleware must maintain both reliability and real-time properties required by the applications despite partial system failure. Traditional fault-tolerance mechanisms usually focus on reliability aspects of the system based on *entity redundancy*. No single point of failure, transparent failover and transparent redirection, and reinvocation are among the requirements of a fault-tolerant ORB (Gokhale et al., 2004). However, less attention has been paid to maintaining desired real-time properties in the presence of faults.

Before describing the fault tolerance techniques in FC-ORB, we first introduce the fault model. FC-ORB is

<sup>2</sup> We note that, as the feedback control loop is designed to control the average utilization within each sampling period, transient deadline misses may occur during a sampling period. Therefore FC-ORB can only provide a statistical guarantee on utilization and deadlines.

designed to handle one or more persistent processor failures. The fault model, known as *fail-stop processors*, is very important and has been considered extensively in the design of fault-tolerant computing systems (Schlichting and Schneider, 1983; Gärtner., 1999). We assume that the communication between the remaining processors does not fail and the network is not overloaded. This is a reasonable assumption for a common class of DRE systems with processor connected with a switched/fast Ethernet LAN with sufficient bandwidth. It should be noted that our utilization control service can be integrated with more sophisticated fault detection and recovery techniques to handle more complex fault models.

FC-ORB improves system robustness in terms of both reliability and real-time properties by integrating three complementary mechanisms. First, FC-ORB provides replication for subtasks and supports transparent failover to backup subtasks located at different processors in face of processor failure. Second, after a processor fails, the remaining processors may experience dramatic workload increase due to the activation of the backup subtasks, which may cause them to miss deadlines or fail. A unique feature of FC-ORB is that it can effectively handle the workload increase via utilization control so that applications can maintain desired real-time properties despite processor failure. Finally, the FC-ORB controller can automatically reconfigure itself at runtime to rebuild its control model, in order to effectively control the DRE system whose deployment is changed due to processor failure.

In our replication mechanism, a subtask may have a backup subtask located on a different processor. For example, the subtask  $T_{13}$  shown in Fig. 1 can have a backup subtask  $T'_{13}$  located on processor  $P_1$ . As a result, when processor  $P_3$  fails because of hardware failure, the execution of subtask  $T_{13}$  is migrated to processor  $P_1$  to continue automatically. Similar to the *COLD\_PASSIVE* replication style used in Fault-Tolerant CORBA (FT-CORBA) (Gokhale et al., 2004), all subtasks are assumed to be stateless (except the connections between subsequent subtasks which are maintained by the middleware) so that the overhead of active state synchronization is avoided.

The failover mechanism works as follows. In the normal mode, each subtask pushes remote operation requests only to the primary instance of its successor. As a result, the backup instance does not receive any requests and its thread remains idle. After a processor fails, the predecessor of a subtask located on the failed processor detects the communication failure based on the underlying socket read/write errors. The predecessor immediately switches the connection to the backup instance of its successor and sends the remote operation requests to it. In the case when the failed processor hosts the first subtask of a task, the controller activates the backup instance of the subtask. Consequently, the execution of the end-to-end tasks is resumed after a transient interruption.

As a part of the fault-tolerant support, the controller in the utilization control service has been designed to be self-

configurable. This is important because the control algorithm relies on knowledge about the subtask allocation in order to compute correct task rates (Lu et al., 2005). When the controller detects communication failure with a processor in the system, it first cancels the periodic timer to pause the feedback control loop. In its internal control model, it then removes the failed processor and moves the subtasks located on the failed processor to the corresponding backup processors. After rebuilding the control model, the controller re-initializes itself and restarts the timer to resume the feedback control loop.

A disadvantage of the centralized control scheme is that the controller becomes a single point of failure. To mitigate this problem, FC-ORB can be easily extended to replicate the controller as well. In the extension, FC-ORB can actively maintain the state consistency between the primary controller and the backup controller, in a way similar to the *ACTIVE* replication style used in FT-CORBA (Gokhale et al., 2004). When the controller executes in replicated mode, all processors send their CPU utilizations to both the primary and the backup controllers at every sampling instant. The backup controller performs control computation just like the primary controller. The difference is that the backup controller does *not* send the resultant new task rates to any processor. Instead, it uses this method to keep the state variables in the backup controller consistent with the primary controller. The primary and backup controllers can exchange heartbeat messages in every sampling period. Once the backup controller stops receiving heartbeats from the primary controller, the backup controller takes over the utilization control service. This feature will allow FC-ORB to maintain control of the entire system even after controller failures.

## 2.5. Implementation

FC-ORB 1.0 has been implemented in C++ using ACE 5.4 on Linux. FC-ORB is based on the FCS/nORB middleware (Lu et al., 2003) which integrates a *single-processor* feedback control scheduling service and a light-weight real-time ORB middleware called nORB (Subramonian et al., 2004). FC-ORB is specialized for memory-constrained DRE systems by supporting a smaller set of features than general-purpose DRE middleware such as TAO. The entire FC-ORB middleware (excluding the code in ACE library and IDL library) is implemented in 7017 lines of C++ code. The controller is implemented in 2089 lines of C++ code and a dynamically linked library that implements the constrained least square solver. We use MATLAB Compiler to create the dynamically linked library from *lsqlin.m* in the MATLAB. At the end of each sampling period, the controller collects the utilizations from application processors and calls the *lsqlin* function in the dynamically linked library with the utilizations as parameters. The *lsqlin* function computes the control input and returns it to the controller. All the code is open-source

and can be downloaded from [http://deuce.doc.wustl.edu/FCS\\_nORB/FC-ORB/](http://deuce.doc.wustl.edu/FCS_nORB/FC-ORB/).

### 3. Empirical evaluation

In this section, we present the results of five sets of experiments run on a distributed testbed with five machines. Experiments I and II evaluate FC-ORB’s performance when task execution times deviate from their estimations and change dynamically at run-time, respectively. Experiment III examines FC-ORB’s capability to handle disturbances from external workloads. Experiment IV tests FC-ORB’s robustness in face of processor failure. Finally, Experiment V measures the overhead introduced by utilization control.

#### 3.1. Experimental setup

All experiments are conducted on a testbed of five machines. All applications and the ORB service run on a Linux cluster composed of four Pentium-IV machines: Ron, Harry, Norbert, and Hermione. Ron and Hermione are 2.80 GHz, and Harry and Norbert are 2.53 GHz. All four machines are equipped with 512 KB cache and 512 MB RAM, and run KURT Linux 2.4.22. The controller is located on another Pentium-IV 2.53 GHz machine with 512 KB cache and 512 MB RAM. The controller machine runs Windows XP Professional. The four machines in the cluster are connected via an internal switch and communicate with the controller machine through the departmental 100 Mbps LAN.

All the experiments run a medium-sized workload that comprises 12 tasks (with a total of 25 subtasks). The tasks include eight end-to-end tasks (tasks  $T_1$  to  $T_8$ ) and four local tasks. Fig. 4 shows how the 12 tasks are distributed on the four application processors. A processor failure incident on Norbert is emulated in Experiment IV to test FC-ORB’s fault-tolerance capability. Hence in Fig. 4, we also show the configured backup subtasks for all subtasks on Norbert that belong to an end-to-end task. There is no

backup subtask for local task  $T_{11,1}$  as we assume that the local task is specific to Norbert.

The subtasks on each processor are scheduled by the RMS algorithm (Liu and Layland, 1973). Each task’s end-to-end deadline is  $d_i = n_i/r_i(k)$ , where  $n_i$  is the number of subtasks in task  $T_i$  and  $r_i(k)$  is the current rate of  $T_i$ . Each end-to-end deadline is evenly divided into subdeadlines for its subtasks. The resultant subdeadline of each subtask  $T_{ij}$  equals its period,  $1/r_i(k)$ . Hence the schedulable utilization bound of RMS (Liu and Layland, 1973),  $B = m(2^{1/m} - 1)$  is used as the utilization set point on a processor, where  $m$  is the number of subtasks (including backup subtasks) on this processor. Specifically, the utilization set points for the four experiment processors are: Ron (72.4%), Harry (72.4%), Norbert (74.3%), and Hermione (72.4%). All (sub)tasks meet their (sub)deadlines if the desired utilization on every processor is enforced. The sampling period of the utilization control service is  $T_s = 4$  s.

To evaluate the robustness of FC-ORB when execution times deviate from the estimations, the execution time of each subtask  $T_{ij}$  can be changed by tuning a parameter called the *execution-time factor*,  $etf_{ij}(k) = a_{ij}(k)/c_{ij}$ , where  $a_{ij}$  is the actual execution time of  $T_{ij}$ . The execution time factor (etf) represents how much the actual execution time of a subtask deviates from the estimation. The etf (and hence the actual execution times) may be kept constant or changed dynamically in a run. In the following we use *inversed etf* (ietf, specifically,  $ietf_{ij}(k) = 1/etf_{ij}(k)$ ) because DRE systems commonly have undesired oscillation when execution times are underestimated (i.e.,  $etf > 1$ ).

We compare FC-ORB against a baseline called OPEN. In OPEN, the utilization control service of FC-ORB is turned off and the middleware becomes a representative real-time ORB without control. OPEN uses a typical open-loop approach to assign task rates based on *estimated* execution time to achieve the desired utilizations. OPEN results in desired utilization when estimated execution times are accurate (i.e.,  $ietf = 1$ ). However, it causes under-utilization when execution times are overestimated (i.e.,  $ietf > 1$ ), and over-utilization when execution times are

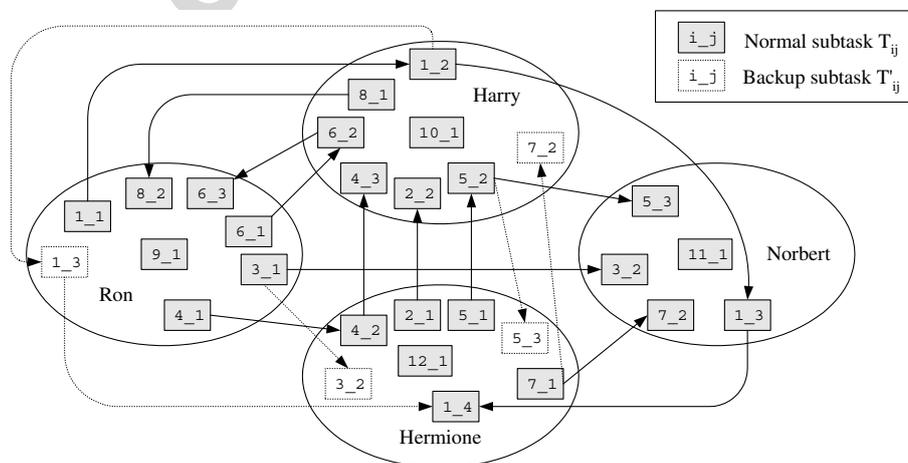


Fig. 4. A medium size workload.

underestimated (i.e.,  $\text{ietf} < 1$ ). This is a common problem faced by application developers because it is often difficult to estimate a tight bound on execution times, especially in unpredictable environments where execution times are heavily influenced by the value of sensor data or user input.

### 3.2. Experiment I: Uncertain execution times

In this subsection, we evaluate FC-ORB's performance when task execution times deviate from the estimations. In each run of this experiment, all subtasks share a fixed execution-time factor ( $\text{ietf}$ ).

First, we run experiments for OPEN which chooses task rates based on estimated execution times so that the estimated utilizations of all processors equal their set points. While the system achieves the desired utilizations in the ideal case when  $\text{ietf} = 1$ , all processors freeze when we set the  $\text{ietf}$  to 0.5. This is not surprising, because the actual execution time of every subtask in the system is *twice* its estimated execution time when  $\text{ietf} = 0.5$ . Consequently, the requested utilization on each processor is about 145% (twice of the desired utilization). Since all FC-ORB threads run at real-time priorities that are higher than the kernel priority on Linux, no kernel activities are able to execute causing the system to crash. This result shows that uncertainties in workloads can significantly degrade the robustness of applications on DRE middleware. On the other hand, the utilizations of all processors drop to only around 18% under OPEN when the actual execution times are only a *quarter* of their estimations ( $\text{ietf} = 4$ ). This results in a extremely underutilized system and unnecessarily low task rates.

In contrast, FC-ORB achieves the desired utilizations on all processors even when execution times deviate significantly from the estimations. Fig. 5(a) shows the utilizations when the average execution time of every subtask is twice its estimation. In the beginning, all processors are overutilized because of the initial task rates. The utilization control service quickly decreases the task rates until the utilizations of all processors converge to the desired levels in around 400 s. Fig. 5(b) shows the utilizations of all processors when the execution time of every subtask is severely overestimated ( $\text{ietf} = 4$ ). In this case, all processors are initialized underutilized due to the low execution times. FC-ORB then increases the task rates until the utilizations of all processors converge to the set points roughly at 500 s. In this experiment, the utilization control service success-

fully prevents the system from crashing and underutilization via rate adaptation.

To examine FC-ORB's performance under different execution time factors, we plot the mean and standard deviation of utilizations of all processors during each run in Fig. 6. Every data point is based on the measured utilization  $u(k)$  from time 1200 s to 1600 s to exclude the transient response at the beginning of each run. FC-ORB consistently achieves the desired utilizations for all tested execution-time factors within the  $\text{ietf}$  range [0.5, 4] which corresponds to eight times variation in execution times. The results show that FC-ORB can enhance system reliability and achieve robust real-time performance under a wide range of operating conditions. Interestingly, when the  $\text{ietf}$  is lower or equal to 0.33, the system freezes due to the extremely high utilization in the beginning of the run. Even though the control thread runs at highest real-time priority, the communication subsystem of Linux runs only at kernel priority. Therefore, the control thread of FC-ORB is blocked on communication because the Linux kernel is preempted by the middleware threads. As a result, the system fails to recover promptly from overload when the  $\text{ietf}$  is equal to or lower than 0.33, even with the help of FC-ORB. In addition, as observed in (Lu et al., 2005), the EUCON algorithm can cause performance oscillation when execution times are underestimated ( $\text{ietf} < 1$ ). Therefore, application developers should use pessimistic estimations of task execution times in FC-ORB. A fundamental advantage of FC-ORB is that it does not cause system underutilization even when task execution times are severely overestimated.

However, we note that some processors fail to reach the utilization set points when  $\text{ietf}$  is equal to or larger than 5. This is because the achievable utilizations are limited by the task rate constraints. For example, when  $\text{ietf}$  is 6, even though the rates of all subtasks on Norbert are adjusted to the maximum values, the utilization of the processor remains below the utilization set point. Note that this is the desired behavior, i.e., task rates are maximized when the system is underloaded.

### 3.3. Experiment II: Varying execution times

The second set of experiments tests FC-ORB's ability to maintain robust real-time performance when task execution times vary *dynamically* at run-time. To investigate

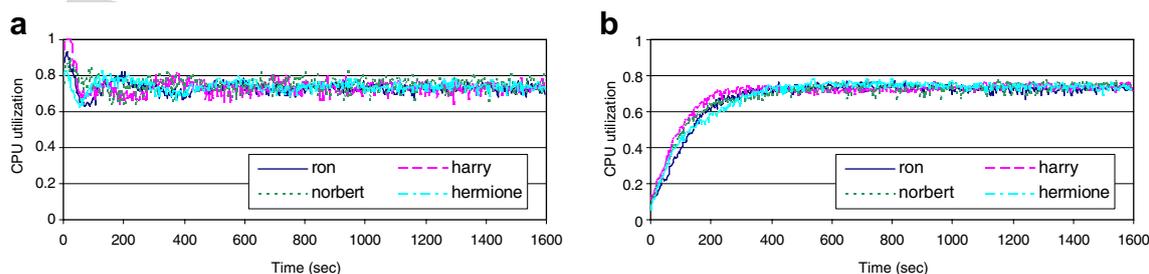


Fig. 5. CPU utilizations under FC-ORB when task execution times deviate from estimations. (a)  $\text{ietf} = 0.5$  (b)  $\text{ietf} = 4$ .

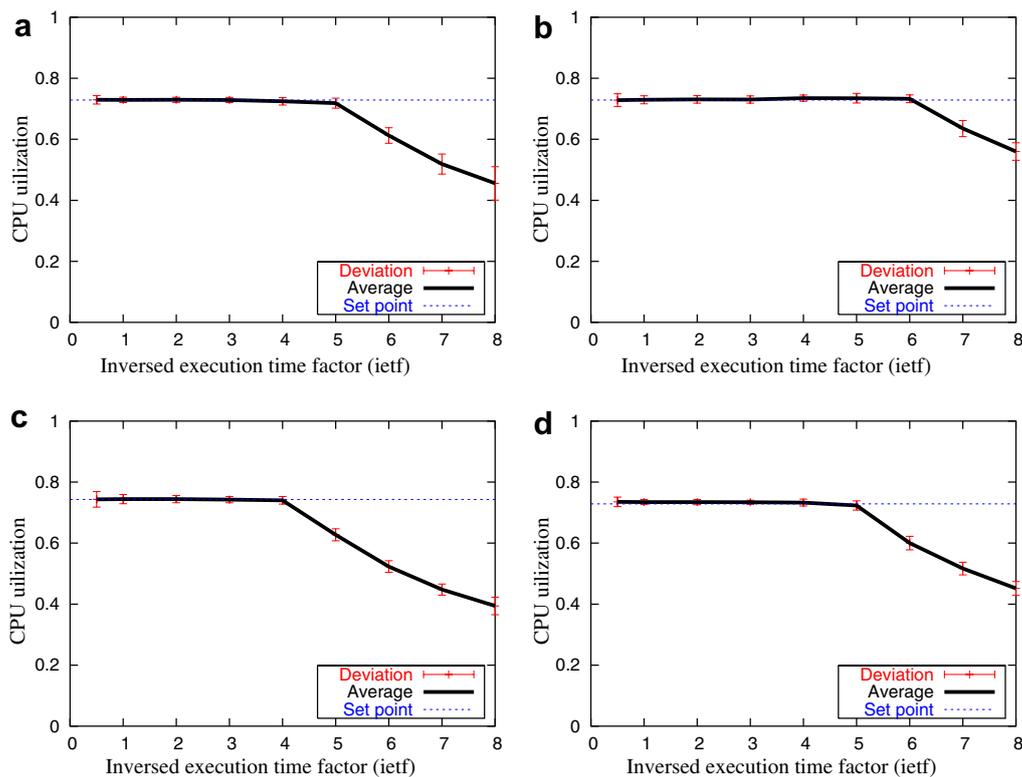


Fig. 6. CPU utilizations of all processors under different execution-time factors. (a) Ron, (b) Harry, (c) Norbert and (d) Hermione.

the robustness of FC-ORB we create two scenarios of workload fluctuation. In the first set of runs, the average execution times on all processors change simultaneously. In the second set of runs, only the execution times on Ron change dynamically, while those on the other processors remain unchanged. The first scenario represents *global* load fluctuation, while the second scenario represents *local* fluctuation on a part of the system.

Fig. 7(a) shows a typical run of OPEN under global workload fluctuation. The ietf is initially 2. At 600 s, it is decreased to 1.33, which corresponds to a 50% increase in the execution times of all subtasks. At time 1000 s, the ietf is increased to 3 to emulate a 56% decrease in execution times. OPEN fails to achieve the desired utilizations due to the lack of dynamic adaptation. In sharp contrast to OPEN, FC-ORB effectively maintains the desired utilizations on all processors under the same workload. As shown in Fig. 7(b), the ietf changes to 1.33 at 600 s such that all processors are suddenly overloaded. FC-ORB responds to the overload condition by decreasing task rates which causes the utilizations on all processors to re-converge to their set points within 100 s (25 control periods). At 1000 s, the utilizations on all processors drop sharply due to the 56% decrease in execution times, causing FC-ORB to dramatically increase task rates until the utilizations re-converge to their set points.

In each run with local workload fluctuation, as shown in Fig. 7(c), the ietf on Ron follows the same variation as the global fluctuation, while all the other processors have a

fixed ietf of 2. As shown in Fig. 7(d), under FC-ORB the utilization of Ron converges to its set point after the significant variation of execution times at 600 s and 1000 s, respectively. We also observe that the other processors experience only slight utilization fluctuation after the execution times change on Ron. This result demonstrates that FC-ORB effectively handles the interdependencies among processors during rate adaptation.

### 3.4. Experiment III: External disturbances

We now evaluate FC-ORB under resource contention from external workloads that are not controlled by FC-ORB. Such external disturbances may be caused by a variety of sources including (i) processing of critical events that must be executed at the cost of other tasks, (ii) varying workload from a different subsystem (e.g., legacy software from a different vendor), and (iii) software faults or adversarial cyber attacks. To stress-test FC-ORB, we emulate the external disturbances using a high priority real-time process to compete with FC-ORB for CPU resource. To investigate the robustness of FC-ORB we create both periodic and aperiodic disturbances. In the first set of runs, the external process *periodically* invokes a function with a *fixed* execution time of 100 ms every 500 ms. In the second set of runs, the external process *aperiodically* invokes another function with a *random* execution time. Both the request interarrival time and the execution time follow exponential

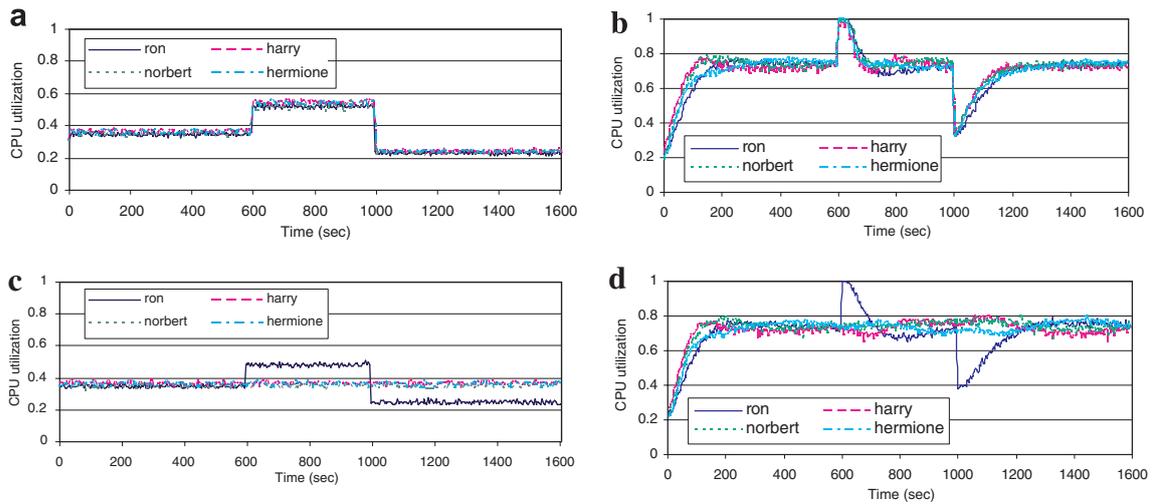


Fig. 7. CPU utilizations of all processors when execution times fluctuate at run-time ( $ietf = 2$ ). (a) OPEN with global fluctuation; (b) FC-ORB with global fluctuation; (c) OPEN with local fluctuation and (d) FC-ORB with local fluctuation.

distributions with mean values of 50 ms and 10 ms, respectively.

The workload controlled by FC-ORB has an  $ietf = 2$ . Here we manually configure the task rates in OPEN such that the workloads achieve the desired utilizations without the external disturbances. As shown in Fig. 8(a), the system does achieve the required performance initially. However, at time 240 s, 360 s, 480 s and 600 s, the external task is activated sequentially on Ron, Harry, Norbert and Hermione. Consequently, the utilizations of all processors are raised to 100%. In contrast to OPEN, Fig. 8(b) shows that FC-ORB successfully maintains the desired utilizations and thus tolerates the external resource contention. Similar situations occur for aperiodic disturbance, except that in this case, both OPEN and FC-ORB have higher fluctuation. Despite noise introduced by the aperiodic requests, FC-ORB still successfully maintains the CPU utilization under

80% most of the time and achieves the desired CPU utilizations on average.

### 3.5. Experiment IV: Processor failure

In this experiment, we evaluate FC-ORB’s ability to recover from processor failure. At 800 s, we emulate the failure of Norbert by using the Linux *kill* command to eliminate the process which carries FC-ORB and the application. The CPU utilization of Norbert immediately drops to almost zero because no other application is running on Norbert. All subtasks on Norbert have backup subtasks located on other processors as shown in Fig. 4, except the local task  $T_{11,1}$ . Their preceding subtasks on other processors detect the communication failure with Norbert and then redirect the remote operation requests to the backup

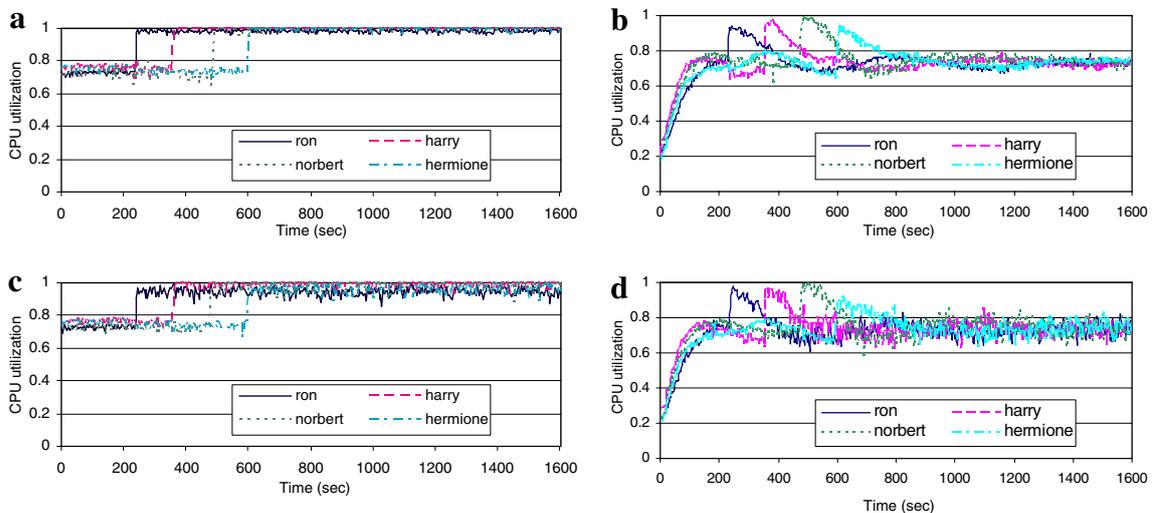


Fig. 8. CPU utilizations of all processors under external disturbances ( $ietf = 2$ ). (a) OPEN with periodic disturbance; (b) FC-ORB with periodic disturbance; (c) OPEN with aperiodic disturbance and (d) FC-ORB with aperiodic disturbance.

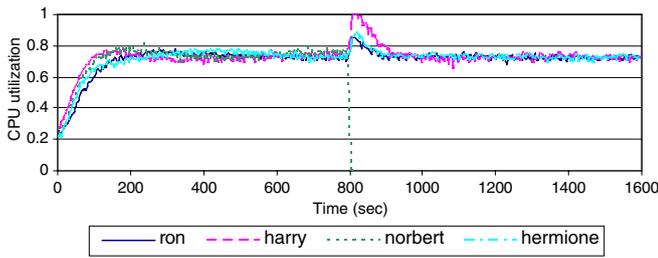


Fig. 9. CPU utilizations of all processors while Norbert has a system failure ( $ietf = 2$ ).

subtasks. Hence, the load of Norbert is distributed to the other three processors in the system.

As demonstrated in Fig. 9, the CPU utilizations of the other three processors increase simultaneously after the failure of Norbert. At the same time, the controller on the control processor re-configures itself to rebuild its control model after it detects the communication failure with Norbert. Thanks to the utilization control service, the high utilizations on the other three processors quickly converge to the desired utilization bounds within 100 s which ensures desired end-to-end real-time performance. Our results demonstrate that the system successfully recovers from a processor failure of a processor and the utilization of the remaining processors converges to a desirable state that ensures the real-time properties of the end-to-end application.

The fault injection using the *kill* command allows us to focus on the robustness of the utilization control service rather than the error detection method. Error detection is a complementary problem to the FC-ORB adaptation for error recovery. Our experimental evaluation of the FC-ORB robustness can be extended to more realistic processor crash failures assuming an appropriate error detection method. The time required for error recovery will include both the time needed for error detection and the convergence of the utilization control service. Formally evaluating the availability of the distributed application requires the definition of an appropriate benchmark (Meyer, 1980), and is a subject of future work.

### 3.6. Experiment V: Overhead

The utilization control service necessarily introduces overhead. This overhead is caused by several factors including the timers associated with FC-ORB, the utilization monitoring, the control computation, the rate enforcement and the thread priority adjustment. Utilization control is a viable middleware service only if the overhead it introduces is sufficiently low. To measure the overhead accurately, we adopt a time stamping approach. Firstly, we differentiate all control service related code from other FC-ORB code. Then, time stamps are taken at the starting point and at the finishing point of each segment of the control service code to get the execution time of the control service. Since the utilization control service runs at the

highest Linux real-time priority, the code segment between two timestamps will not be preempted during its execution. Hence, the time-stamped result accurately reflects the real execution overhead.

To achieve fine grained measurements, we adopt a nano-second scale time measuring function called *gethrtime*. This function uses an OS-specific high-resolution timer that returns the number of clock cycles since the CPU was powered up or reset. The *gethrtime* function has a low overhead and is based on a 64 bit clock cycle counter on Pentium processors. With the clock counter number divided by the CPU speed, we can get reasonably precise and accurate time measurements.

Table 1 lists the average and standard deviation of the overhead of the utilization monitor, the actuator (including the rate modulator and the priority adjuster) and the controller of the utilization control service. All results in the table are obtained from over 600 continuous sampling periods. The overhead of the utilization monitor is very low because it just executes around 20 lines of code to read the utilization data from the Linux system file */proc/stat*.

The actuator has the dominant overhead because it involves relatively more complicated operations. The rate modulator and the priority manager are the two main contributors to the actuating overhead. Our implementation uses the ACE function *reset\_timer\_interval* to reset the timers and the ACE function *thr\_setprio* to adjust the thread priorities in FC-ORB. In most cases, only the rate modulator is invoked to adapt the task rates by adjusting the interval of the timers. In some periods when the order of the task rates has been reversed, the priority manager is invoked to adjust the priorities of the real-time threads. The overhead of adjusting thread priorities is much larger than resetting timer intervals and so the standard deviation of the actuating overhead is large.

To estimate the average computation overhead of the controller, we measure the execution time of the *lsqlin* function in the shared library which dominates the computation cost on the control processor. We call the *lsqlin* function for 1000 times as a subroutine. The result is then divided by 1000 to get the execution time of a single execution of the least square computation. As shown in Table 1, the overhead of the controller is stable with small deviation and its amount is between that of the monitor and the actuator. Overall, the execution time overhead of all control

Table 1  
Overhead of utilization control

Processor	Monitor (ms)		Actuator (ms)		Controller (ms)	
	Avg.	Dev.	Avg.	Dev.	Avg.	Dev.
Ron	0.090	0.013	19.078	18.160		
Harry	0.096	0.013	34.389	33.305		
Norbert	0.094	0.012	39.460	37.223		
Hermione	0.088	0.013	27.924	25.951		
Controller					5.765	0.219

components in our experiments is around 46 ms per sampling period, corresponding to 1.15% utilization given a sampling period of 4 s.

#### 4. Related work

Adaptive middleware is emerging as a core building block for DRE systems. For example, TAO (Schmidt, 2002), dynamicTAO (Kon et al., 2002), ZEN (Klefstad et al., 2002), and nORB (Subramonian et al., 2004) are adaptive middleware frameworks that can (re)configure various properties of ORB middleware at design- and run-time. Higher-level adaptive resource management frameworks, such as QuO (Zinky et al., 1997), Kokyu (Gill et al., 2001) and RT-ARM, leverage lower-level mechanisms provided by ORB middleware to (re)configure scheduling, dispatching, and other QoS mechanisms in higher-level middleware. ORB services such as the TAO Real-Time Event Service (Harrison et al., 1997) and TAO Scheduling Service (Gill et al., 2001) offer high-level services for managing reliability and real-time properties of interactions between application components. FC-ORB has several important features that distinguishes itself from earlier work on adaptive middleware. First, FC-ORB integrates the end-to-end scheduling service with a utilization control service. This integrated approach enables the middleware to meet end-to-end deadlines by dynamically controlling the utilizations on individual processors. Second, in contrast to earlier works that rely on heuristics-based adaptive techniques, FC-ORB implements control algorithms that has been rigorously designed and analyzed based on a control-theoretic approach. Finally, FC-ORB enhances traditional fault-tolerance mechanisms with utilization control techniques to handle processor failures.

Several other projects also applied control theoretic approaches to real-time systems. For example, Steere et al. (1999), developed a feedback based CPU scheduler that coordinated allocation of CPU cycles to consumer and supplier threads in a modified Linux kernel. Abeni et al. (2002) presented analysis of a reservation-based feedback scheduler. In Lu et al. (2002), a set of feedback control real-time scheduling algorithms were proposed to provide deadline miss ratio and utilization guarantees for single-processor systems. Feedback control real-time scheduling has also been extended to handle distributed systems with independent tasks (Stankovic et al., 2001). For systems requiring discrete control adaptation strategies, hybrid control theory has been adopted to control state transitions among different system configurations (Abdelwahed et al., 2003; Koutsoukos et al., 2005). A key difference between the work presented in this paper and the related work is that we describe the design and evaluation of a utilization control service in an ORB middleware, while the related work is based either on simulations or kernel implementations. ORB middleware is a particularly suitable layer for managing *end-to-end* adapta-

tion in distributed systems since it operates at a broader (distributed) scope than stand-alone operating systems.

In our earlier work we studied EUCON (Lu et al., 2005) only through control-theoretic analysis and simulation results. FC-ORB implements and empirically evaluates the end-to-end utilization service on an ORB middleware and a physical testbed. Furthermore, we also extend the EUCON algorithm with controller reconfiguration and replication techniques for handling processor failures.

Agilos (Li and Nahrstedt, 1999) and ControlWare (Zhang et al., 2002) were two earlier control-based middleware framework for QoS adaptation. They are targeted at multimedia and Internet servers instead of DRE applications. FCS/nORB (Lu et al., 2003) is another real-time ORB middleware that features a feedback control real-time scheduling service. However, FCS/nORB only controls the real-time performance of a single server in a client/server environment. In contrast, FC-ORB provides an *end-to-end* utilization control service in a peer-to-peer architecture for DRE systems. A key feature of FC-ORB is that it can effectively coordinate the adaptation on multiple interdependent processors through a distributed feedback control loop.

#### 5. Conclusions

In summary, we have designed and implemented FC-ORB, a real-time ORB middleware with a novel end-to-end utilization control service. Our experiments on a physical testbed has shown that (1) FC-ORB can enforce desired utilizations on all processors in a DRE system, even when task execution times deviate significantly from their estimated values or vary significantly at run-time; (2) FC-ORB can survive considerable resource contention imposed by external disturbances; (3) FC-ORB enhances the robustness of real-time properties to processor failures; (4) the middleware layer instantiation of the end-to-end utilization control service only introduces a small amount of processing and memory overhead. These results demonstrate that the integration of end-to-end utilization control, fault-tolerance mechanisms, and end-to-end scheduling in ORB middleware is a promising approach to achieve robust real-time performance guarantees for DRE applications. In the future, we plan to enhance FC-ORB to incorporate other adaptation mechanisms such as admission control and task reallocation so that FC-ORB can be applied to a broader class of applications. An important research direction is to integrate FC-ORB with advanced fault detection and recovery in order to handle more complex fault models.

#### References

- Abdelwahed, S., Neema, S., Loyall, J.P., Shapiro, R., 2003. A hybrid control design for QoS management. In: Proc. of RTSS.
- Abeni, L., Palopoli, L., Lipari, G., Walpole, J., 2002. Analysis of a reservation-based feedback scheduler. In: Proc. of IEEE RTSS, December.

- Brandt, S., Nutt, G., Berk, T., Mankovich, J., 1998. A dynamic quality of service middleware agent for mediating application resource usage. In: Proc. of IEEE RTSS, December.
- Gärtner, F.C., 1999. Fundamentals of fault-tolerant distributed computing in asynchronous environments. *ACM Comput. Surv.* 31 (1), 1–26.
- Gill, C.D., Levine, D.L., Schmidt, D.C., 2001. The design and performance of a real-time CORBA scheduling service. *Real-Time Syst., The Int. J. Time-Critical Comput. Syst.* 20 (2), Special issue on Real-Time Middleware.
- Gokhale, A.S., Natarajan, B., Schmidt, D.C., Cross, J.K., 2004. Towards real-time fault-tolerant CORBA middleware. *Cluster Comput.* 7 (4), 331–346.
- Harrison, T., Levine, D., Schmidt, D., 1997. The Design and Performance of a Real-time CORBA Event Service. In: Proc. of OOPSLA.
- Kao, B., Garcia-Molina, H., 1997. Deadline assignment in a distributed soft real-time system. *IEEE Trans. Parallel Distrib. Syst.* 8 (12), 1268–1274.
- Klefsstad, R., Schmidt, D.C., O’Ryan, C., 2002. Towards highly configurable real-time object request brokers. In: Proc. of Symp. on Object-Oriented Real-Time Distributed Computing, pp. 437–447.
- Kon, F., Costa, F., Blair, G., Campbell, R.H., 2002. The case for reflective middleware. *Commun. ACM* 45 (6), 33–38.
- Koutsoukos, X., Tekumalla, R., Natarajan, B., Lu, C., 2005. Hybrid supervisory utilization control of real-time systems. In: Proc. of IEEE RTAS.
- Lehoczky, J.P., 1990. Fixed priority scheduling of periodic task sets with arbitrary deadline. In: Proc. of IEEE RTSS.
- Li, B., Nahrstedt, K., 1999. A control-based middleware framework for QoS adaptations. *IEEE J. Selected Areas Commun.* 17 (9), 1632–1650.
- Liu, J.W.S., 2000. *Real-Time Systems*. Prentice Hall.
- Liu, C., Layland, J., 1973. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM* 20 (1), 46–61.
- Lu, C., Stankovic, J., Tao, G., Son, S., 2002. Feedback control real-time scheduling: Framework, modeling, and algorithms. *Real-Time Syst.* 23 (1/2), 85–126.
- Lu, C., Wang, X., Gill, C., 2003. Feedback control real-time scheduling in ORB middleware. In: Proc. of IEEE RTAS, May.
- Lu, C., Wang, X., Koutsoukos, X., 2005. Feedback utilization control in distributed real-time systems with end-to-end tasks. *IEEE Trans. Parallel Distrib. Syst.* 16 (6), 550–561.
- Marti, P., Fohler, G., Fuertes, P., Ramamritham, K., 2002. Improving quality-of-control using flexible timing constraints: metric and scheduling. In: Proc. of IEEE RTSS.
- Meyer, J.-F., 1980. On evaluating the performability of degradable computing systems. *IEEE Trans. Comput.* 29 (8), 720–731.
- Natale M.D., Stankovic, J., 1994. Dynamic end-to-end guarantees in distributed real-time systems. In: Proc. of IEEE RTSS.
- Schlichting, R.D., Schneider, F.B., 1983. Fail-stop processors: An approach to designing fault-tolerant computing systems. *ACM Trans. Comput. Syst.* 1 (3), 222–238.
- Schmidt, D., 1995. Reactor: An object behavioral pattern for concurrent event demultiplexing and event handler dispatching. In: Coplien, J.O., Schmidt, D.C. (Eds.), *Pattern Languages of Program Design*. Addison-Wesley, Reading, MA, pp. 529–545.
- Schmidt, D., 1997. Acceptor and connector: Design patterns for initializing communication services. In: Martin, R., Buschmann, F., Riehle, D. (Eds.), *Pattern Languages of Program Design*. Addison-Wesley, Reading, MA.
- Schmidt, D.C. et al., 2002. TAO: A Pattern-Oriented Object Request Broker for Distributed Real-time and Embedded Systems. *IEEE Distributed Syst. Online* 3 (2).
- Schmidt, D.C., Kuhns, F., 2000. An overview of the real-time CORBA specification. *IEEE Comput.* 33 (6), 56–63.
- Seto, D., Lehoczky, J.P., Sha, L., Shin, K.G., 1996. On task schedulability in real-time control system. In: Proc. of IEEE RTSS.
- Stankovic, J.A., et al., 2001. Feedback control scheduling in distributed systems. In: Proc. of IEEE RTSS.
- Steere, D.C., Goel, A., Gruenberg, J., McNamee, D., Pu, C., Walpole, J., 1999. A feedback-driven proportion allocator for real-rate scheduling. *Operating Syst. Des. Implementation*, 145–158.
- Subramonian, V., Xing, G., Gill, C.D., Lu, C., Cytron, R., 2004. Middleware specialization for memory-constrained networked embedded systems. In: Proc. of RTAS.
- Sun, J., Liu, J.W.-S., 1996. Synchronization protocols in distributed real-time systems. In: Proc. of ICDCS.
- Wang, X., Lu, C., Koutsoukos, X., 2005. Enhancing the robustness of distributed real-time middleware via end-to-end utilization control. In: Proc. of IEEE RTSS.
- Zhang, R., Lu, C., Abdelzaher, T.F., Stankovic, J.A., 2002. ControlWare: A Middleware Architecture for Feedback Control of Software Performance. In: Proc. of ICDCS.
- Zhao, F., Koutsoukos, X.D., Haussecker, H.W., Reich, J., Cheung, P., Picardi, C., 2001. Distributed monitoring of hybrid systems: A model-directed approach. In: Proc. of IJCAI, pp. 557–564.
- Zinky, J.A., Bakken, D.E., Schantz, R.E., 1997. Architectural support for quality of service for CORBA objects. *Theor. Practice Object Syst.* 3 (1).