

CReMeS: A CORBA Compliant Reflective Memory based Real-time Communication Service

Sun-Tae Chung, Oscar Gonzalez, Krithi Ramamritham, Chia Shen

TR2000-18 May 2000

Abstract

We present CReMeS, a CORBA-compliant design and implementation of a new real-time communication service. It provides for efficient, predictable, and scalable communication between information producers and consumers. The CReMeS architecture is based on MidART's Real-Time Channel-based Reflective Memory (RT-CRM) abstraction. This architecture supports the separation of QoS specification between producer and consumer of data and employs a user-level scheduling scheme for communicating real-time tasks. These help us achieve end-to-end predictability and allows our service to scale. The CReMeS architecture provides a CORBA interface to applications and demands no changes to the ORB layer and the language mapping layer. Thus, it can run on non real-time Off-The-Shelf ORBs and enables applications on these ORBs to have scalable and end-to-end predictable asynchronous communication facility. In addition, an application designer can select whether to use an out-of-band channel or the ORB GIOP/IOP for data communication. This permits a trade-off between performance, predictability and reliability. Experimental results demonstrate that our architecture can achieve better performance and predictability %but perhaps a reduced level of reliability than a real-time implementation of the CORBA Event Service when the out-of-band channel is employed for data communication; it delivers better predictability with comparable performance when the ORB GIOP/IOP is used.

IEEE 2000 Real-Time Systems Symposium, November 27-30, 2000. Orlando, Florida, USA.

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

CReMeS: A CORBA Compliant Reflective Memory based Real-time Communication Service

Sun-Tae Chung

Oscar González
Chia Shen

Krithi Ramamritham

TR-2000-18 November 2000

Abstract

We present CReMeS, a CORBA-compliant design and implementation of a new real-time communication service. It provides for efficient, predictable, and scalable communication between information producers and consumers. The CReMeS architecture is based on MidART's Real-Time Channel-based Reflective Memory (RT-CRM) abstraction. This architecture supports the separation of QoS specification between producer and consumer of data and employs a user-level scheduling scheme for communicating real-time tasks. These help us achieve end-to-end predictability and allows our service to scale. The CReMeS architecture provides a CORBA interface to applications and demands no changes to the ORB layer and the language mapping layer. Thus, it can run on non real-time Off-The-Shelf ORBs and enables applications on these ORBs to have scalable and end-to-end predictable asynchronous communication facility. In addition, an application designer can select whether to use an out-of-band channel or the ORB GIOP/IIOP for data communication. This permits a trade-off between performance, predictability and reliability. Experimental results demonstrate that our architecture can achieve better performance and predictability than a real-time implementation of the CORBA Event Service when the out-of-band channel is employed for data communication; it delivers better predictability with comparable performance when the ORB GIOP/IIOP is used.

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Information Technology Center America; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Information Technology Center America. All rights reserved.

In Proceedings of the 21st IEEE Real-Time Systems Symposium, November 27-30, 2000. Orlando, Florida, USA



CReMeS: A CORBA Compliant Reflective Memory based Real-time Communication Service*

Sun-Tae Chung[†]

Oscar González[‡]

Krithi Ramamritham[‡]

Chia Shen[§]

Abstract

We present CReMeS, a CORBA-compliant design and implementation of a new real-time communication service. The CReMeS architecture is based on MidART's Real-Time Channel-based Reflective Memory (RT-CRM) which (a) supports the separation of QoS specification between producer and consumer of data and (b) employs a user-level scheduling scheme for communicating real-time tasks. Although RT-CRM provides for predictable and scalable asynchronous communication infrastructure for applications, it is not a standard middleware. The desire to make its service widely available motivated us to develop CReMeS by adapting RT-CRM for CORBA environments. As a result, CReMeS provides for efficient, predictable, and scalable communication between information producers and consumers. Since CReMeS provides a CORBA interface to applications and demands no changes to the ORB layer and the language mapping layer, it can run on non real-time Off-The-Shelf ORBs and enables applications on these ORBs to have scalable and end-to-end predictable asynchronous communication facility. In addition, an application designer can select whether to use an out-of-band channel or the ORB GIOP/IIOP for data communication. This permits a trade-off between performance, predictability and reliability. Experimental results demonstrate that our architecture can achieve better performance and predictability than a real-time implementation of the CORBA Event Service when the out-of-band channel is employed for data communication; it delivers better predictability with comparable performance when the ORB GIOP/IIOP is used.

1 Introduction

The need to reduce the time and effort required to deploy and maintain distributed real-time applica-

tions has necessitated a move towards using off-the-shelf (OTS) distributed object computing middleware based on open standards. The Common Object Request Broker Architecture (CORBA) [10] specified by Object Management Group (OMG) is beginning to gain acceptance as such a standard middleware platform for use even in distributed real-time applications. However, when using CORBA as part of a distributed real-time architecture problems arise due to its communication primitives being based on method invocations and from its lack of support for expressing and handling timing constraints [19, 20, 23]. To provide support for QoS requirements of real-time applications, especially those that need communication beyond method invocations, a *Real-Time Event Service* [7, 17] has been designed by extending CORBA's *event channel* abstraction [11]. The CORBA Event Service model was originally intended for event dispatching and thus it is only appropriate for applications where asynchronous exchange of small size data message is all that is required. Also since all asynchronous communication traffic is supposed to go through the event channel, the event channel can become a bottleneck. However, if this bottleneck is avoided by using multiple event channels, the real-time scheduler associated with one event channel may not be aware of messages being exchanged by other event channels and their resource requirements and hence admission control becomes a problem. Therefore, the RT Event Service alone may have difficulty in providing satisfactory end-to-end QoS. Furthermore, a large class of distributed real-time applications, such as process control, factory automation, and telecommunications [22], require predictable communication services capable of handling both *small (e.g., remote operation invocations and event notification)* and large messages (e.g., audio/video). A new CORBA service that allows an application to use the same programming model for different data types is not only desirable, but it can also reduce the complexity of developing these types of applications.

To address the above deficiencies, we propose CReMeS (a CORBA Compliant **R**eflective **M**emory based Real-Time Communication Service):

*Research supported in part by NSF Grant CDA-9502639, by Mitsubishi Electric Research Labs, and by Soongsil University.

[†]School of Electronic Engineering - Soongsil University - Seoul, 156-743, Korea - cst@syscon.soongsil.ac.kr

[‡]Computer Science Dept. - Univ. of Massachusetts - Amherst, Mass. 01003, USA - {ogonzale, krithi}@cs.umass.edu

[§]MERL(CRL) - Mitsubishi Electric Research Labs (Cambridge Research Lab), Cambridge, Mass. 02139, USA - shen@merl.com

- CReMeS' *efficiency* derives from its use the Real-Time Channel-based Reflective Memory (RT-CRM) abstraction [21] used in MidART [6] as the architectural basis for communication between nodes.
- CReMeS provides for *end-to-end predictability* by utilizing MidART's admission control at both the supplier and consumer nodes, and the user-level scheduling scheme for communicating distributed real-time tasks [22].
- CReMeS is *scalable* since the scheduler in one node handles the communication traffic relating only to that node, as opposed to the scheduler in the Real-Time Event Service which handles all the traffic connecting to the event channel.
- CReMeS' *flexibility* arises from the fact that application designers can
 - (a) separately specify the QoS requirements for suppliers and consumers, thus allowing different timing properties to be met independently.
 - (b) use a CORBA interface to applications, thus obviating any need for changes to the ORB layer and the language mapping layer. Thus, it can run on non real-time OTS (Off-The-Shelf) ORBs and enables applications on these ORBs to have scalable and end-to-end predictable asynchronous communication facility.
 - (c) use an out-of-band channel or the ORB GIOP (General Inter-ORB Protocol) and IIOP (Internet Inter-ORB Protocol) for data communication. This permits a trade-off between performance and reliability.

There are two ways of incorporating new functionality into CORBA. The first consists of integrating the functionality into the ORB core and/or the language mapping layer. But this method loses portability (because the implementation is ORB-dependent) and interoperability (because both clients and servers have to use the same ORB implementation). The second is a Service approach [4] which does not require modification of the ORB core and language mapping layer. In this service approach, the module implementing the new functionality is defined by only an IDL interface and thus is independent of the ORB and language mapping. The latter, service-based, approach is CORBA compliant and hence is desirable.

While a CORBA compliant adaptation of MidART's RT-CRM is desirable – since it is portable and interoperable, and thus can run on any ORBs – it is not trivial to achieve. One of the major challenges comes from the need to achieve object location transparency while still keeping the efficient shared memory based mechanism that is optimal for intra-node communication.

Other challenges include the design of interfaces and adaptation of communication paths of MidART while not losing the performance advantages of MidART.

Keeping the goals of achieving performance, predictability, and scalability in mind, we have designed and implemented a prototype of the CReMeS architecture using OmniORB2 version 3.0.0 [16], an open source ORB, on Windows NT 4.0. Experimental results demonstrate that our architecture can achieve better performance (as measured by latency), predictability (as measured by jitter and the ability to provide QoS guarantees) and scalability (measured in terms of the number of clients or size of messages handled) than a real-time implementation of the CORBA Event Service when an out-of-band channel is employed. It displays superior predictability and scalability while delivering comparable performance when the standard two-way synchronous communication of the ORB GIOP/IIOP is used for data streams. This makes our solution adapt to the needed reliability and performance requirements for building performance-sensitive real-time applications in CORBA environments.

This paper is organized as follows: Section 2 provides a brief description of CORBA, and a critique of its real-time capability. In Section 3, we present a description of RT-CRM in the context of MidART. The design and implementation details of CReMeS is presented in Section 4. The experimental performance results of CReMeS implementation are described in Section 5. Related work is reviewed in Section 6. We concluded the paper in Section 7.

2 Overview of CORBA

CORBA, a middleware standard specified by OMG (Object Management Group)[10], defines how objects in heterogeneous distributed environments can be described and can interact with each other. The main component of CORBA is the Object Request Broker (ORB), basically a software bus responsible for locating objects and delivering clients' requests to server objects. An ORB provides object location and object implementation transparency in addition to communication infrastructure. CORBA Services, another component of CORBA, are general-purpose high level distributed services on top of ORB to provide convenient functionality to applications. The interface to CORBA Services is defined by OMG. Adopted OMG Object Services, collectively called CORBAServices (COS) include the naming service – which allows clients to find objects based on names, and the event service – which supports asynchronous events (data) delivery through appropriate event channel implementations. In order to support a general ORB interoperability, CORBA specifies GIOP which is the implementation of GIOP over TCP/IP.

“Request/Response” type of operation invocations supported by CORBA are usually implemented by a synchronous two-way communication model where clients synchronously wait for response from a server once they invoke operations on server objects. While this type of synchronous two-way communication model is useful in general, for industrial control and multimedia applications timely delivery of message and continuous media data delivery are necessary and so the time taken by a client to wait for server response may cause unacceptable delays.

For these reasons, CORBA’s Event Service was defined to provide asynchronous communication facility for CORBA applications by decoupling between suppliers of (event) data and consumers of (event) data. CORBA’s Event Service can provide anonymous, and one-to-many or many-to-many communication among suppliers and consumers and make connection configuration possible at run-time. The recently adopted CORBA Messaging Specification [12] introduced the type-safe Asynchronous Method Invocation (AMI) model and is aimed at one-to-one communication but is not flexible enough to be configured at run-time even though AMI invocations can specify QoS requirements.

2.1 Use of CORBA in Real-Time Environments

Current CORBA specifications lack several real-time features such as the standard interfaces through which real-time applications can specify their QoS requirements to the ORB, and configure and control ORB resources appropriately for the achievement of desired real-time QoS. Some of these deficiencies have been alleviated with the recent introduction of CORBA Messaging QoS policy specification [12] and Real-Time CORBA specification [13]. These specifications are aimed at soft real-time applications. There is still no standard way for clients to indicate timing requirements like latency and jitter of their requests. Furthermore, compared to low-level programming approaches, current ORB implementations incur significant run-time overheads affecting both throughput and latency, and additional unpredictability [19]. These weaknesses have been studied extensively [19, 20, 23] and have led to the development of real-time ORB support. TAO [19] is a high performance ORB (TAO) whose key components include a real-time I/O system, real-time inter-ORB protocol engine, real-time object adapter, real-time scheduling and dispatching mechanism, an optimized IDL compiler which generates efficient and predictable stubs and skeletons. An alternative approach of Wolfe et. al. [23] involves extending the current ORB (as opposed to designing a new Real-Time ORB like TAO) to deal with real-time needs. To this end, they

have developed a Real-Time manager and object services such as global time service, real-time event service, global priority service, and real-time scheduling to equip current ORB with real-time capability. It should be noted that even these RT-ORBs cannot get rid of all of unpredictability because they don’t have complete control of the resources that they need.

An alternative to above extensions to the ORB is to build a higher level CORBA service to provide required QoS to applications. One such service, is the RT Event Service [7, 17], developed to extend COS Event Service into real-time application domain. RT Event Service extends CORBA Event Service by supporting periodic rate-based event processing and efficient event filtering and correlation. Even though the Real-Time Event Service provides the above additional features and good performance under light load conditions, it is not scalable nor predictable under heavy load conditions. These drawbacks of the Event Service are mainly due to fundamental architectural coupling of each event channel with all of its associated suppliers and consumers. Our proposed CReMeS addresses these shortcomings.

3 Real-Time Channel-based Reflective Memory (RT-CRM)

Since CReMeS is based on MidART’s Real-Time Channel-based Reflective Memory (RT-CRM), in this section, we discuss the relevant aspects of RT-CRM. RT-CRM was proposed as a service in the MidART middleware. Figure 1 shows the high level architecture of RT-CRM. RT-CRM is a software-based reflective memory – it provides data reflection with guaranteed timeliness. Data reflection is defined as the memory-to-memory data transfer among application host memories in a networked environment.

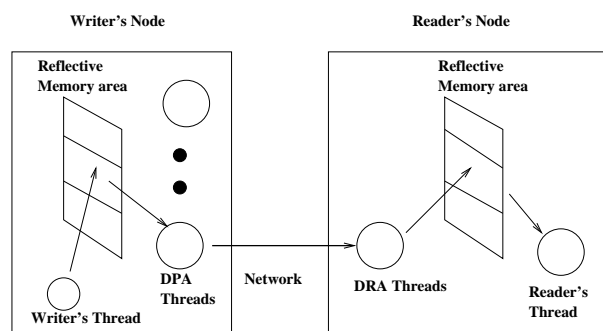


Figure 1. RT-CRM High Level Architecture

Data reflection is accomplished by a Data Push Agent (DPA) residing on the writer’s node and sharing the writer’s memory area. This agent represents the reader’s QoS and data reflection requirements. A virtual channel is established between the agent and the

reader’s memory area, through which the writer’s data is actively transmitted and written into the reader’s local memory area by a Data Receive Agent (DRA). RT-CRM provides efficient and flexible asynchronous communication. The run-time communication model of a channel-based reflective memory is defined by the data sending and data reception semantics. Data sending can be either Push-on-Write or Periodic Push, while data reception can be either blocking or non-blocking. Since DPA and DRA are separate threads of control from the application threads, the data sending and data reception semantics are entirely definable and parameterized by the subscriber. The combination of the data push and reception semantics lead to several general models of communications commonly found in distributed real-time applications [22].

4 CReMeS Architecture for CORBA environments

CReMeS implements MidART’s RT-CRM as a service for the CORBA environment. In this section, we describe the principal design issues and our implementation experience of CReMeS.

4.1 Design considerations and options

We had three design goals:

1. Maintaining compatibility with MidART’s current usage.
2. Avoiding any performance loss compared to using MidART without CORBA.
3. Providing the application with CORBA development semantics.

We now elaborate on how we achieve these design goals in CReMeS.

CReMeS retains compatibility with MidART’s original API and their semantics. In MidART, the interface to the RT-CRM service is provided through a library of API similar to memory read-write syntax. In CReMeS, the interface to the RT-CRM service is constructed as a CORBA object which is defined by IDL and exposed to applications. This CORBA object is called the RT-CRM Interface Object (RIO). Inside the RIO servant (i.e., the entity implementing the CORBA object [10]), a mapping to the reflective memory is provided. RIO methods have the same invocation semantics as MidART’s original API. Thus by invoking methods of RIO, an application can set up and use a RT-CRM service. The RIO interface allows applications to specify real-time requirements in terms of deadlines, reader periods, writer periods, and guarantee/non-guarantee mode, and to choose Push-on-Write or Periodic Push mode of communication. This interface is intuitive and much simpler to use than other CORBA communication services such as the RT Event Service.

The procedures to use a RT-CRM service in CORBA environments is similar to those in MidART explained in [25]. The only exception is that an application (writer or reader) now needs to obtain the object reference to RIO before the application can actually invoke any of the methods.

As for the second goal, there are many design considerations, but two important ones are the location of RIO and the communication path between the reflective memories.

The location of RIO affects the performance of the communication service. There are two options for the location. One is to locate RIO in an address space different from that of the application, and the other is to locate the object in the same address space as that of the application.

If RIO is located in an address space different from the application which manages the reflective memory, one natural choice of path for data stream transfer between an application (writer or reader) and the associated reflective memory is the ORB GIOP/IIOP path. In this design choice, even asynchronous one-way data transfer needs to pass through the ORB GIOP/IIOP path twice – once at the writer’s end and again at the reader’s end – before reaching its final destination. This not only incurs overhead but also increases the jitter. From our initial experimental results obtained with this design choice, the potential for poor performance (increased latency and high jitter) was obvious, especially when transmitting data of large size. In order to achieve performance similar to that of MidART, we need to adopt the shared memory mechanism employed in MidART for data stream transfer between the writer/reader and the reflective memory that it is attached to. In MidART, through the shared memory mechanism, a writer (or reader) can ask for data to be copied directly into (from) the reflective memory in the same node without additional overheads (such as going through IP loopback or additional memory copies). However, if the servant is located in an address space different from the application, it is difficult to make use of the shared memory mechanism, especially if we want to keep the same interface semantics as MidART’s library API.

Thus, the best choice for the location of the RIO is the application’s address space. Inside the servant of the RIO, a shared memory mechanism for data stream transfer between an application and reflective memory is provided. This data transfer path does not incur much of the ORB GIOP/IIOP overheads since in most of currently existing OTS ORBs (including omniORB2 [16]), object invocation on an object in the same address space (called “collocated object”) is optimized to be a virtual function call.

Let us now examine the issue of the communication path between reflective memories. Since the ORB provides a communication infrastructure, one can replace every communication path (including data reflection between reflective memories) in the implementation of RT-CRM by the ORB GIOP/IOP path. While the ORB GIOP/IOP provides reliable communication, it incurs overhead at both the GIOP/IOP layer and TCP layer which GIOP/IOP is based on. From our experimental results, we found that in order to achieve performance similar to that of MidART but under CReMeS we need to use a separate data channel for data reflection between two reflective memories which does not make use of the ORB GIOP/IOP path. To this end, we make use of a separate UDP channel for this path. It should be pointed out that we employ the ORB GIOP/IOP for setting up a RT-CRM service since a reliable communication path is desirable in this case.

This approach of adopting separate communication paths in CORBA environments is also seen in CORBA audio/video service specification [14]. There, for the configuration and control of streams, it uses the GIOP/IOP of the ORB. Once the streams are configured for an audio/video service, a different communication channel, such as UDP, is used for data stream transmission. However, there exists a trade-off between gaining performance by adoption of a separate UDP channel and losing reliability by not using the ORB IOP.

In order to achieve the third design goal, the design consideration that we faced can be described as follows. In CORBA environments, the location of the target CORBA object is supposed to be transparent to a client. A client usually contacts a naming server and obtains the object references to the target CORBA object. How can we achieve location transparency of the RIO if the (target) CORBA object is located in the same address space as the client? Our solution to this conundrum is to place the (target) RIO in the (dynamic) library to be linked with the application. When the library is initialized, the servant of the RIO is instantiated inside the library, registered into the ORB, and the object reference to the CORBA object is created and registered into a naming server. The library to be linked with the application at compile-time has information about the servant of RIO, but the location information is hidden from the application. An application is given the IDL of the interface at compile-time (and also the library to be linked), and then it contacts the naming server at run-time and obtains the object reference of the CORBA object. One copy of RIO is created for each application process and registered with a unique name. The uniqueness of the name is achieved by using process identifier as a part of the name.

Thus, applications using CReMeS have the same usage semantics as typical CORBA applications.

4.2 Components of CReMeS

CReMeS is composed of three basic modules: the RT-CRM Server, a Library and the RT-CRM Service Coordinator.

The *RT-CRM Server* is responsible for managing reflective memories, scheduling messages and dispatching them to the other RT-CRM Servers. It contains a scheduler, a dispatcher, and an admission control module. The scheduler employs dual priority user-level scheduling scheme [2] in order to integrate non real-time data as well as real-time data by providing fair bandwidth for non real-time data in such a way that real-time communication is not affected. The detailed structures of scheduler, dispatcher, and admission control are the same as those in MidART (for details, the readers should refer to [22]). The RT-CRM Server exposes two CORBA objects, “Lib_Comm” and “CRM_Comm”. Through the interface of “Lib_Comm” CORBA object, the library can deliver setup messages to the RT-CRM Server. “CRM_Comm” CORBA object provides an interface to the RT-CRM Server through which the other RT-CRM Servers can deliver setup messages when necessary. The RT-CRM server runs as a process in a node, one copy per node.

The *library* supports mapping to RT-CRM for applications (producers or consumers). An important component of the library is the servant of RIO discussed in the previous subsection. The library has the necessary support to allow RIO to setup the shared memory area between an application and the RT-CRM Server. Also, the library is responsible for locating and calling the interface of the CORBA object “Lib_Comm” that allows the transfer of data between an application and the RT-CRM Server.

RT-CRM Service Coordinator provides a set of administrative services including registration, and location of RT-CRM services for the entire system. It contains the necessary information that allows the RT-CRM Server to locate and establish connections between the reader and writer of a reflective memory. It is implemented as a single process and exposes one CORBA object (“Glob_Comm”) to RT-CRM Servers.

4.3 CReMeS Communication Structure

In CReMeS, three different kinds of communication paths are supported.

1. A path to setup a RT-CRM service: All communications for setup pass through the ORB GIOP/IOP path. The three modules (the library, the RT-CRM Server and the RT-CRM Coordinator) use this communication path to set up a RT-CRM service.

2. A path for data stream communication between an application (producer or consumer) and the associated reflective memory (in the same node): Producers write to their reflective memory through the shared memory mechanism. Similarly, making use of the shared memory mechanism, consumers read from the reflective memories in the RT-CRM Server in the consumer's node.
3. A path for data stream communication between producer's reflective memory and consumer's reflective memory: Data streams written into the reflective memory are scheduled and dispatched by the scheduler and dispatcher within the RT-CRM Server, and are copied or "reflected" into the reflective memory in the RT-CRM Server in the consumer nodes through a UDP channel (or the ORB GIOP/IIOP path).

5 Experiments and Performance Evaluation

We have done extensive experimentation to evaluate the real-time properties of CReMeS and its ability to integrate real-time and non-real-time tasks. We begin with an overview of the experimental setup and a description of the metrics used in the evaluation, and then turn to the performance results.

5.1 Experimental Hardware and Software

All of the experiments were carried out on Pentium II PCs with Windows NT 4.0 operating systems over 100BaseT Fast Ethernet. We have isolated the network segment from the rest of the LAN while running all the experiments. The two PCs we used in the experiments have processor speeds of 333MHz and 266MHz respectively. The faster machine is equipped with 128MB of RAM and the slower machine has 64MB of RAM.

We ran the experiments described in this section on two real-time communication services, CReMeS and the Real-Time Event Service v1.9 which is included in the distribution of TAO v1.1.1.

The current implementation of CReMeS uses a robust, high-performance ORB, omniORB2 v3.0.0, developed by AT&T [16]. This ORB implements the specification 2.3 of the Common Object request Broker Architecture (CORBA). It supports the C++ language binding, is fully multi-threaded and comes with a COS Naming Service. We evaluated the CReMeS architecture under its two data communication mechanisms: one using an out-of-band channel (UDP) and the other using ORB GIOP/IIOP.

TAO's Real-Time Event Service is an object-oriented, real-time implementation of the CORBA Event Service intended to decouple suppliers and consumers and allow asynchronous event delivery in a predictable manner. RT Event Service runs on top of TAO, a real-time ORB

endsystem that provides end-to-end quality of service guarantees [7]. It is important to note that TAO and its implementation of the Event Service can be configured in a variety of ways to optimize performance. The results presented in this section were obtained for the optimized version of TAO and the RT-Event Service.

5.2 Metrics and Experimental Setup

We evaluate the benefits of our real-time communication service as well as TAO's RT Event Service by measuring their *performance* and *predictability* as a function of load. For this purpose, we measure the *latency* and *jitter* introduced by the communication service under different load conditions. In addition, we also determine how well the communication service *scales* with respect to the sizes of messages using the service.

We used a simple distributed application that we believe is indicative of the type of real-time load that is likely to be present in a distributed industrial control system [9]. In this application, a producer of sensor data (representing a data acquisition device, such as a remote terminal unit or programmable logic controller) gathers data from sensors at regular intervals and sends it to an operator workstation; the operator workstation processes the data arriving from the sensors and displays the results on the screen. In all of our experiments, this consuming process does not perform any processing of the data.

The producer and the consumer processes ran on two nodes connected via an isolated network. In one node, the *producer* application can be configured to run 1 to N threads, each thread representing a different device. In the other node, the *consumer* application creates 1 to N entities to receive the messages. Each *producer* thread, P_0 to P_N , sends messages through the communication service to its corresponding *consumer* entity, C_0 to C_N , in the other node. In our experiments, we vary the size of message, the production rate of messages and the number of producer/consumer pairs to evaluate the real-time properties of the communication channel in question.

5.3 Latency and Jitter Measurements

High latency overheads introduced by the communication service have a negative impact in meeting deadlines. However, high variance in the latency has a more detrimental effect for real-time applications that require end-to-end predictability.

In this experiment the producer creates n threads for each run. *Producer* thread P_0 sends a message to its corresponding consumer C_0 every 50 msec and it is assigned the highest priority. The other producers, P_1 to P_n , send a message to their corresponding consumers, C_1 to C_n , every 100 msec and are assigned a lower priority than P_0 . Both producer and consumer first create




Figure 2. Latency for High Priority thread, 1KB




Figure 3. Latency for High Priority thread, 8KB

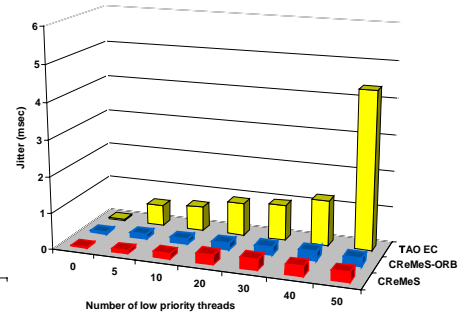


Figure 4. Jitter for High Priority thread, 8KB

their application threads and set up the necessary associations by contacting the communication service. All threads block on a signal until all of the associations between producers and consumers have been accepted and established by the communication service. The highest priority thread, P_0 , is responsible for releasing all other threads. Once the signal to run is given, threads are dispatched according to their priority and the communication service is responsible for the scheduling of messages between the two nodes.

Each *producer* of data includes in its message a timestamp. Upon receiving the message, a *consumer* returns this timestamp to *producer* via a separate UDP channel in order to determine the round trip latency of each message. *Jitter* is determined by calculating the standard deviation of the latency for a particular run. In each run the highest priority producer sends 4000 messages and the lower priority producers send 2000 messages. We vary the load in the experimental set up by increasing the number of low priority producer/consumer pairs from 0 to 50 and the size of the messages from 1KB to 8KB. Each message is a sequence of values of type CORBA::Octet.

Figures 2 and 3 show the average latency produced by the high priority producer/consumer pair for each of the communication services evaluated (CReMeS using the out band channel, CReMeS-ORB using omniORB2 and TAO RT-Event Service) for message sizes of 1KB and 8KB respectively. The jitter results for the high priority producer/consumer pair with 8KB messages for all the communications services are shown in Figure 4. Due to space limitations, we do not present the jitter results for messages of size 1KB. However, it is important to note that the jitter observed for all three communications services under all load conditions is less than 0.2 msec for the high priority thread and less than 1 msec for the low priority threads.

Figures 5 and 6 show the average latency produced by one of the low priority producer/consumer pairs for

each of the communication services for message sizes of 1KB and 8KB respectively. The jitter results for low priority thread with 8KB size messages are shown in Figure 7.

Overall, CReMeS performs better than the rest of the services tested:

- CReMeS has a lower latency, under most load conditions, for both the high priority and the low priority producer/consumer pairs than the other services.
- Perhaps, more importantly the variance in latency for CReMeS is quite low and remains stable as the load increases.

As Figure 2 shows, for CReMeS, latency values for the high priority producer range from 0.88 msec with no low priority producers to 0.99 msec (a 12% increase) with 50 low priority producers for 1KB. It ranges from 1.8 to 2.6 msec (a 42% increase) for messages of size 8KB.

- For all load conditions the jitter observed for the high priority thread when using CReMeS is less than the one observed using the RT-Event Service and similar to CReMeS-ORB.

Figure 4 shows that jitter for the high priority producer/consumer pair ranges from 0.025 to 0.3 msec for 8KB messages.

- The jitter for the low priority producer/consumer pair is higher than the one recorded for the high priority pair. However, as Figure 7 shows, the jitter variations for 8KB messages for the low priority pair are smaller for CReMeS than the ones produced by other services.

The low variation in jitter for both high priority and low priority producer/consumers pairs observed with CReMeS for different loads is an indication of the high level of *predictability* and *scalability* that CReMeS is capable of providing. A predictable and scalable service is useful to an application designer, since it allows him/her




Figure 5. Latency for Low Priority thread, 1KB




Figure 6. Latency for Low Priority thread, 8KB




Figure 7. Jitter for Low Priority thread, 8KB

to determine a realistic bound for the worst case execution time incurred by the communication service. Figure 8 and Figure 9 clearly illustrate the above point.

Figure 8 shows a trace of the latency measurements taken for the high priority producer/consumer pair when 40 low priority producers are active and sending messages of 8KB. Figure 9 shows a similar trace for the same load, but instead of using CReMeS as the communication service TAO’s RT-Event Channel is used. Figure 10 presents the average latency for the high priority thread with 40 low priority threads for each communication as a function of the size of the message. Similarly, in Figure 11 the average latency for one of the 40 low priority threads for the same load, 40 threads, is shown for the same messages sizes. Clearly, CReMeS is the only communication service that scales, in terms of message size, for both high and low priority producers.

Turning to the specifics of the performance of the other threads, TAO’s RT Event Service results indicate that this service is capable of providing predictable performance only for small messages or a small number of producer/consumer pairs. Figure 2 shows that the latency for the high priority producer increases 20%, from 1.9 to 2.3 msec, for 1KB size messages. An increase in latency of 93%, from 3.1 msec to 6.1 msec, is observed in Figure 3 for the high priority producer/consumer pair with 8KB messages as the number of low priority producers increases from 0 to 50. The percentage of increase in latency for both messages sizes is higher than the increase recorded for both CReMeS and CReMeS-ORB. Similarly, the latency for the low priority producer shows a similar increase of 17% and 128% for 1KB and 8KB message sizes respectively.

The jitter results for the RT Event Service also show a similar trend, however the increase is more pronounced for the high priority producer/consumer pair, as shown in Figure 4. The jitter for the low priority producer/consumer pair, shown in Figure 7, increases considerably once the number of low priority

producer/consumer pairs is greater than 10, but remains constant, ranging from 1.8 msec to 2 msec from then on.

These results indicate that TAO RT Event Service scales poorly on our test platform. However, our experimental results show that the RT Event Service performance is acceptable when the size of the message is low and the amount of traffic passing through the event channel is moderate.

CReMeS-ORB provides similar latency and jitter as CReMeS for the high priority producer/consumer pair as CReMeS and similar to TAO’s RT Event Channel for the low priority pair. Overall, CReMeS-ORB consistently provides lower latency than the RT Event Channel for most of the loads evaluated, as shown in Figures 2, 5, 3 and 6. However, the results presented in Figure 7, indicate that this communication service introduces more jitter for the low priority producer/consumer than the other two services for most of the load conditions. One possible reason for this is that CReMeS-ORB uses the two-way synchronous request, instead of the one-way operations used by RT Event Service. As with the RT Event Service, this variation in jitter once the load increases does not bode well for systems that require predictable QoS.

6 Related Work

ARMADA[1] is a middleware supporting fault-tolerance and end-to-end guarantee for embedded real-time distributed applications, but unlike CReMeS, is not constructed on a standard interface. Rajkumar et.al [18] propose the Real-Time publisher/subscriber model. This model is similar to event service and the RT-CRM communication paradigm used in CReMeS in that they all provide anonymous and group communication. But one key difference between the RT-CRM communication paradigm and many of the current publisher-subscriber communications models is that RT-CRM effectively achieves the decoupling of the writer/publisher’s quality of service (QoS) characteris-




Figure 8. Latency trace CReMeS




Figure 9. Latency trace for RT-Event Service

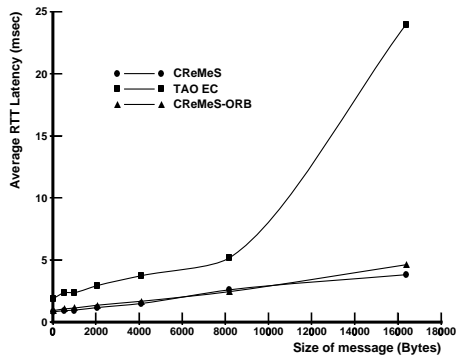


Figure 10. Latency for high priority thread




Figure 11. Latency for low priority thread

tics from the reader/subscriber’s QoS requirement. This decoupling allows much more flexibility in constructing distributed and/or concurrent real-time applications. What this means is that RT-CRM is not simply a multicast or unicast data transport protocol. It is not restricted to transmit the data to the receivers/readers immediately after the data is made available by a writer. In contrast, it provides application designers the facility to specify how and when the data should be sent according to the application’s specific needs.

Recently, OMG adopted a Notification Service [15] which extends the COS Event Service by adding event filtering, and event delivery QoS to event service functionalities. However, this Notification Service will have the same limitation as COS Event Service since it employs the same event channel architecture as COS Event Service does.

In [8], TMOSM (TMO Support Middleware) is developed as a TMO(Time-triggered Message-triggered Object) execution engine in CORBA environments. TMOSM is adapted from original TMO execution engine to be used in CORBA environments just as our CReMeS is adapted from original MidART. Finally, our work has goals similar to that of the DARPA Quorum program and associated projects [5, 24] which are aimed at enabling or enforcing QoS capabilities for CORBA.

7 Conclusions

In this paper, we have described our experience with the design, implementation, and evaluation of CReMeS, a CORBA compliant communication service based on the concept of a Real-Time Channel-based Reflective Memory abstraction. CReMeS was designed as an alternative communication mechanism to the Event Service in order to provide adequate end-to-end QoS.

We have designed and implemented this service and have tested our implementation under various load conditions. Our experimental results demonstrate that our architecture can achieve better performance (measured by latency), predictability (measured by jitter) and scalability (measured in terms of number of clients or size of messages handled) than a real-time implementation of the CORBA Event Service when an out-band-channel is employed. These results are important for applications that required asynchronous communication of large size data. It is to be noted that these results were obtained even though CReMeS was implemented on top of an ORB which has not been optimized for real-time applications. If CReMeS ran on top of real-time ORB and/or real-time operating systems, we expect the predictability of the communication services to improve even further.

We designed a variation of CReMeS to use the ORB

as the communication channel for the reflective memories, instead of UDP. Our results indicate that CReMeS-ORB consistently provides lower latency than the RT Event Service. We plan to conduct further experiments and modifications in this architecture to improve its performance. For some applications this variation may be desirable since it enhances the reliability of the communication because of ORB's use of TCP.

As part of our future work, we will focus on the dispatching mechanism for the RT-CRM Server module. Currently, CReMeS employs a single dispatching thread. In [22], we have shown that a single dispatcher thread can minimize priority inversion, enable priority tracking and helps in dealing with limited operating system priority levels. However, this mechanism works well when the sending thread is not blocked after passing a message to the network layer. This is not the case in our current implementation of CReMeS-ORB, which may account for the difference in performance between CReMeS-ORB and the RT Event Service under heavy load conditions. One of the future direction is to include a multi-threaded dispatching mechanism in CReMeS-ORB that works in unison with our current implementation of a dual priority user level scheduler.

References

- [1] T. Abdelzaher et al., "ARMADA Middleware and Communication Services," *Real-Time Systems*, vol. 16, no. 2/3, pp. 127-153, 1999, Kluwer Academic Publishers.
- [2] R. Davis and A. Wellings, "Dual Priority Scheduling," *Proc. of 16th IEEE Real-Time Systems Symposium*, December, 1995.
- [3] DARPA:1999, "The Quorum Program", www.darpa.mil/ito/research/quorum/index.html.
- [4] P. Felber, B. Garbinato, and R. Guerraoui, "The Design of a CORBA Group Communication Service," *Proc. of the 15th IEEE Symposium on Reliable and Distributed Systems (SRDS'96)*, October, 1996.
- [5] W. Feng, U. Syyid, and J.S. Liu, "Providing for an Open, Real-time CORBA," *Proc. of the Workshop on Middleware for Real-time systems and Services*, San Francisco, CA, December, 1997.
- [6] O. González, C. Shen, "Implementation and Performance of MidART". *Proc. of IEEE Workshop on Middleware for Distributed Real-Time Systems and Services*, San Francisco, CA., December, 1997.
- [7] T. Harrison, L. Levine and D. C. Schmidt, "The Design and Performance of a Real-time CORBA Event Service," *Proc. of OOPSLA '97*, ACM, Atlanta, GA, October, 1997.
- [8] K.H. Kim, M. Ishida, and J. Liu, "An Efficient Middleware Architecture Supporting Time-Triggered Message-Triggered Objects and an NT-based Implementation," *Proc. of IEEE CS Int'l Symp. On Object-Oriented Real-time Distributed Computing*, 1999.
- [9] K. Ramamritham, C. Shen, O. González, S. Sen and S. Shirgurkar, "Using Windows NT for Real-Time Applications: Experimental Observations and Recommendations," *Proc. of 4th IEEE Real-Time Technology and Applications*, Denver, Colorado, June, 1998.
- [10] The Common Object Request Broker: Architecture and Specification, *OMG Document 99-10*.
- [11] CORBA Services: Common Object Services Specification, *OMG Document 98-12*.
- [12] CORBA Messaging Specification, *OMG TC Document orbos/98-05-05*.
- [13] Real-Time CORBA Joint Revised Submission, *OMG Document ptc/99-05-03*.
- [14] Control and Management of Audio/Video Streams, *OMG RFP Submission, OMG TC Document telcom/97-05-07*.
- [15] CORBA Notification Service, *OMG TC Document telcom/99-07-01*.
- [16] omniORB2, <http://www.uk.research.att.com/omniORB>.
- [17] C. O'Ryan, Douglas C. Schmidt, D. Levine, and R. Noseworthy, "Applying a Real-time CORBA Event Service to Large-scale Distributed Interactive Simulation," *Proc. IEEE of 5th International Workshop on Object-oriented Real-Time Dependable Systems (WORDS '99)*, Monterey, CA., November, 1999.
- [18] R. Rajkumar, M. Gagliardi, and L. Sha, "The Real-Time Publisher/Subscriber Interprocess Communication Model for Distributed Real-Time Systems: Design and Implementation," *Proc. of 1st IEEE Real-Time Technology and Application Symposium*, May, 1995.
- [19] D. C. Schmidt, D. Levine, and S. Mungee, "The Design of the TAO Real-Time Object Request Broker," *Computer Communications Special Issue on Building Quality of Service into Distributed Systems*, Elsevier Science, Volume 21, No. 4, April, 1998.
- [20] D. C. Schmidt, S. Mungee, S. Flores-Gaitan, and A. Gokhale, "Software Architectures for Reducing Priority Inversion and Non-determinism in Real-time Object Request Brokers," to appear, *Real-Time Systems*, 2000, Kluwer Academic Publishers.
- [21] C. Shen and I. Mizunuma, "RT-CRM: Real-Time Channel-based Reflective Memory," *Proc. of 3rd IEEE Real-Time Technology and Applications Symposium*, Montreal, Canada, June, 1997.
- [22] C. Shen, O. González, K. Ramamritham and I. Mizunuma, "User Level Scheduling of Communicating Real-Time Tasks," *Proc. of 5th IEEE Real-Time Technology and Applications Symposium*, Vancouver, Canada, June, 1999.
- [23] V. Wolfe, L. Dipippo, R. Ginis, M. Squadrito, S. Wohlever, I. Zyk and R. Johnston, "Expressing and Enforcing Timing Constraints on a Dynamic Real-Time CORBA System," *Real-Time Systems*, vol. 16, 253-280, 1999, Kluwer Academic Publishers.
- [24] J.A. Zinky, D.E. Bakken, and R. Schantz, "Architectural Support for Quality of Service for CORBA Objects," *Theory and Practice for Object Systems*, Vol. 3, no.1, 1997.
- [25] MidART User Guide, <http://www.merl.com/projects/midart/index.html>