

# Evaluating Policies and Mechanisms to Support Distributed Real-Time Applications with CORBA

Carlos O’Ryan and Douglas C. Schmidt

{coryan,schmidt}@uci.edu

Electrical & Computer Engineering Dept.

University of California, Irvine

Irvine, CA 92697, USA

This paper appeared in the proceedings of the Sixth IEEE Real-Time Technology and Applications Symposium (RTAS’00), Washington D.C., USA, 31 May – 2 June 2000.

Fred Kuhns, Marina Spivak, Jeff Parsons,

Irfan Pyarali, and David L. Levine

{fredk,marina,parsons,irfan,levine}@cs.wustl.edu

Department of Computer Science, Washington University

St. Louis, MO 63130, USA\*

**Subject Areas:** Real-time CORBA; Patterns and Frameworks; Distributed and Real-Time Middleware

## Abstract

*To be an effective platform for performance-sensitive real-time systems, commodity-off-the-shelf (COTS) distributed object computing (DOC) middleware must support application quality of service (QoS) requirements end-to-end. However, conventional COTS DOC middleware does not provide this support, which makes it unsuited for applications with stringent latency, determinism, and priority preservation requirements. It is essential, therefore, to develop standards-based, COTS DOC middleware that permits the specification, allocation, and enforcement of application QoS requirements end-to-end.*

*The Real-time CORBA and Messaging specifications in the CORBA 2.4 standard are important steps towards defining standards-based, COTS DOC middleware that can deliver end-to-end QoS support at multiple levels in distributed and embedded real-time systems. These specifications still lack sufficient detail, however, to portably configure and control processor, communication, and memory resources for applications with stringent QoS requirements.*

*This paper provides four contributions to research on real-time DOC middleware. First, we illustrate how the CORBA 2.4 Real-time and Messaging specifications provide a starting point to address the needs of an important class of applications with stringent real-time requirements. Second, we illustrate how the CORBA 2.4 specifications are not sufficient to solve all the issues within this application domain. Third, we describe how we have implemented portions of these specifications, as well as several enhancements, using TAO, which is our open-source real-time CORBA ORB. Finally, we evaluate the performance of TAO empirically to illustrate how its features address the QoS requirements for certain classes of real-time applications.*

---

\*This work was funded in part by AFOSR grant F49620-00-1-0330, Boeing, NSF grant NCR-9628218, DARPA contract 9701516, Motorola, Nortel, SAIC, Siemens, and Sprint.

## 1 Introduction

**Challenges for next-generation real-time systems:** Due to the need to handle stringent constraints on efficiency, predictability, memory footprint, and weight/power consumption, software techniques used to develop real-time systems have historically lagged behind those used to develop mainstream desktop and server software. As a result, real-time software applications are difficult to evolve and maintain. Moreover, they are often so specialized that it is not cost effective to adapt them to leverage new technology innovations or to meet new market opportunities.

To exacerbate matters, a growing class of distributed real-time systems require end-to-end support for various quality of service (QoS) aspects, such as latency, jitter, and throughput. These applications include the control and management of telecommunication systems, commercial and military aerospace systems, and streaming audio/video over the Internet. In addition to requiring support for stringent QoS requirements, these types of systems are often targeted for markets where deregulation, global competition, and/or R&D budget constraints necessitate increased software productivity.

Requirements for increased software productivity motivate the use of distributed object computing (DOC) *middleware* [1], such as CORBA [2] and Java RMI [3]. DOC middleware resides between applications and the underlying operating systems, protocol stacks, and hardware in complex distributed and embedded real-time systems. The *technical goal* of DOC middleware is to simplify software development by shielding applications from component location, programming language, OS platform, communication protocols and interconnects, and hardware dependencies [4]. The *business goal* of DOC middleware is to decrease the cycle-time and effort required to develop real-time applications and services.

In theory, middleware can simplify the creation, composition, and configuration of real-time applications without incurring significant time and space overhead. In practice, however, technical challenges have impeded the development and

deployment of efficient, predictable, and scalable middleware for real-time systems. In particular, commodity-off-the-shelf (COTS) DOC middleware generally lacks (1) support for QoS specification and enforcement, (2) integration with high-speed networking technology, and (3) efficiency, predictability, and scalability optimizations [5]. These omissions have limited the rate at which performance-sensitive applications, such as teleconferencing and avionics mission computing, have been able to leverage advances in DOC middleware.

**Candidate solution → CORBA:** First-generation DOC middleware was not targeted for high-performance and real-time systems. Thus, it was not appropriate for systems with stringent deterministic and statistical real-time QoS requirements [5]. Over the past two years, however, the use of CORBA-based DOC middleware for real-time applications has increased significantly in aerospace [6], telecommunications [7], medical systems [8], and distributed interactive simulation [9] domains. The increased adoption of CORBA stems from the following factors:

1. *The maturation of patterns* – In recent years, a substantial amount of R&D effort has focused on patterns [10, 11]. For instance, research conducted as part of the DARPA Quorum project [12, 7, 5, 13] has identified key patterns [14] and optimization principles [15] for high-performance and real-time systems.
2. *The maturation of frameworks* – Recent progress in patterns R&D has enabled the creation of higher-quality frameworks [16], such as ACE [17], that support the development of QoS-enabled DOC middleware and applications.
3. *The maturation of standards* – During the past decade, the OMG’s suite of CORBA standards has matured considerably, particularly the Real-time [18] and Messaging [19] specifications that define components and QoS features for high-performance and real-time systems.
4. *The maturation of COTS CORBA products* – An increasing number of COTS ORBs [20] are applying patterns and frameworks to implement the CORBA Real-time [5] and Messaging [21] specifications.

The vehicle for our research on DOC middleware for high-performance and real-time applications is TAO [5]. TAO is an open-source CORBA-compliant ORB designed to support applications with stringent end-to-end QoS requirements. In our prior work on TAO, we have shown that it is possible to achieve high efficiency, predictability, and scalability in ORB middleware by applying appropriate concurrency [14], connection [22] and demultiplexing [15] patterns [1].

Our earlier work, however, has not addressed techniques for balancing competing real-time application requirements for la-

tency and throughput. Moreover, the OMG has recently approved the Real-time [18] and Messaging [19] specifications, which give application developers greater control over end-to-end priority preservation and ORB predictability. Therefore, in this paper, we evaluate these specifications to illustrate the extent to which they do and do not satisfy the requirements of an important class of real-time applications. For situations where CORBA 2.4 is under-specified, we demonstrate how the specification can be enhanced to allow greater application control and portability.

**Paper organization:** The remainder of this paper is organized as follows: Section 2 presents an overview of the OMG CORBA specifications relevant to this paper; Section 3 presents standard CORBA 2.4 features as well as TAO extensions, and describes how they address the needs of an avionics mission computing system, a representative of a class of hard real-time applications;

Section 4 evaluates the results of benchmarks that measure efficiency, predictability, and scalability of key CORBA real-time and messaging features in TAO; and Section 5 presents concluding remarks.

## 2 Synopsis of CORBA 2.3 and 2.4 Features

This section describes the CORBA reference model and highlights the difference between CORBA 2.3 (and earlier CORBA specifications) and CORBA 2.4, focusing on features pertaining to quality of service (QoS).

### 2.1 The CORBA 2.3 Reference Model

CORBA Object Request Brokers (ORBs) allow clients to invoke operations on distributed objects without concern for object location, programming language, OS platform, communication protocols and interconnects, and hardware [4]. Figure 1 illustrates the key components in the CORBA reference model upto and including CORBA 2.3 [23] that collaborate to provide this degree of portability, interoperability, and transparency.<sup>1</sup> Each component in the CORBA reference model is outlined below:

**Client:** An application plays the client *role* if it obtains references to objects and invokes operations on them to perform application tasks. Objects can be remote or collocated relative to the client. Ideally, a client can access a remote object just like a local object, *i.e.*, `object→operation(args)`. Figure 1 shows how the underlying ORB components described below

<sup>1</sup>This overview only focuses on the CORBA components relevant to this paper. For a complete synopsis of CORBA’s components see [2].

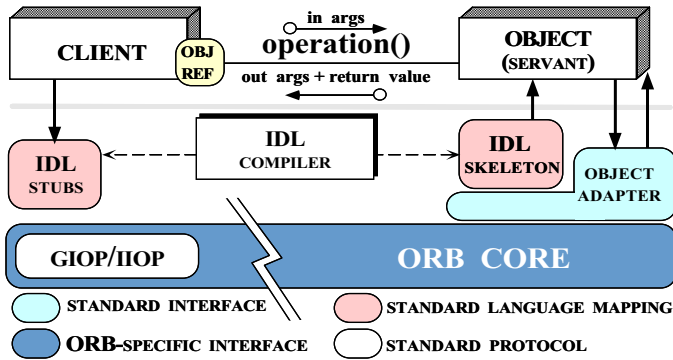


Figure 1: Key Components in the CORBA 2.3 Reference Model

transmit remote operation requests transparently from client to object. Applications can play both the client and server roles.

**Object:** In CORBA, an object is an instance of an OMG Interface Definition Language (IDL) interface. Each object is identified by an *object reference*, which associates one or more paths through which a client can access an object on a server. An *object ID* associates an object with its implementation, called a servant. Over its lifetime, an object has one or more servants associated with it that implement its interface.

**Servant:** This component implements the operations defined by an OMG IDL interface. In object-oriented (OO) languages, such as C++ and Java, servants are implemented using one or more class instances. In non-OO languages, such as C, servants are typically implemented using functions and structs. A client never interacts with servants directly, but always through objects identified by object references.

**ORB Core:** When a client invokes an operation on an object, the ORB Core is responsible for delivering the request to the object and returning a response, if any, to the client. An ORB Core is usually implemented as a run-time library linked into client and server applications. For objects executing remotely, a CORBA interoperability compliant ORB Core communicates via a version of the General Inter-ORB Protocol (GIOP), such as the Internet Inter-ORB Protocol (IIOP) that runs atop the TCP transport protocol. In addition, custom Environment-Specific Inter-ORB protocols (ESIOPs) can also be defined.

**OMG IDL Stubs and Skeletons:** IDL stubs and skeletons serve as a “glue” between the client and servants, respectively, and the ORB. Stubs implement the *Proxy* pattern [10] and provide a strongly-typed, *static invocation interface* (SII) that marshals application parameters into a common message-level representation. Conversely, skeletons implement the *Adapter* pattern [10] and demarshal the message-level representation

back into typed parameters that are meaningful to an application.

**IDL Compiler:** An IDL compiler transforms OMG IDL definitions into stubs and skeletons that are generated automatically in an application programming language, such as C++ or Java. In addition to providing programming language transparency, IDL compilers eliminate common sources of network programming errors and provide opportunities for automated compiler optimizations [24].

**Object Adapter:** An Object Adapter is a composite component that associates servants with objects, creates object references, demultiplexes incoming requests to servants, and collaborates with the IDL skeleton to dispatch the appropriate operation upcall on a servant. Object Adapters enable ORBs to support various types of servants that possess similar requirements. This design results in a smaller and simpler ORB that can support a wide range of object granularities, lifetimes, policies, implementation styles, and other properties.

## 2.2 QoS-related Enhancements to CORBA 2.4

CORBA specifications upto and including CORBA 2.3 [23] lacked features that allow applications to allocate, schedule, and control key CPU, memory, and networking resources necessary to ensure end-to-end quality of service. The CORBA 2.4 standard [2] includes the Messaging [19] and Real-time CORBA specifications [18] that support many of these features. The Messaging specification defines asynchronous operation models [21] and a QoS framework that allows applications to control many end-to-end ORB policies. The Real-time CORBA specification defines interfaces and policies for managing ORB processing, communication, and memory resources. Figure 2 illustrates how these various CORBA 2.4 features interact.

As shown in Figure 2 an ORB endsystem [5] consists of network interfaces, operating system I/O subsystems and communication protocols, and CORBA-compliant middleware components and services. The CORBA 2.4 specification identifies capabilities that must be *vertically* (*i.e.*, network interface ↔ application layer) and *horizontally* (*i.e.*, peer-to-peer) integrated and managed by ORB endsystems to ensure end-to-end predictable behavior for *activities*<sup>2</sup> that flow between CORBA clients and servers.

Below, we outline these capabilities, starting from the lowest level abstraction and building up to higher-level services and applications.

<sup>2</sup>An activity represents the end-to-end flow of information between a client and its server that includes the request when it is in memory, within the transport, as well as one or more threads.

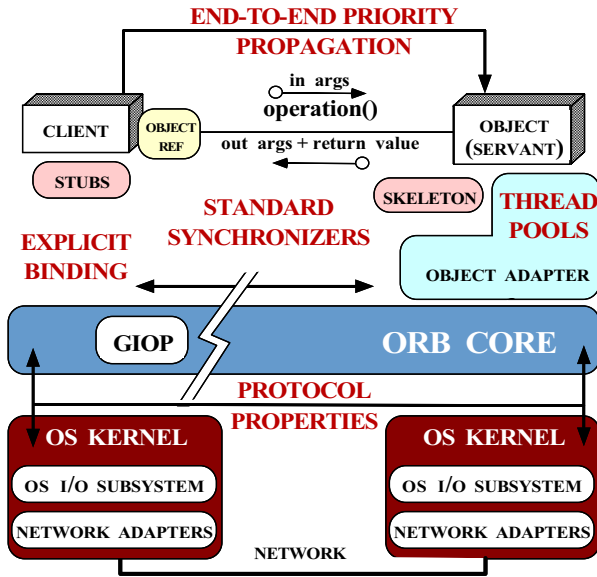


Figure 2: CORBA 2.4 QoS Support for Real-Time Applications

### 1. Communication infrastructure resource management:

A CORBA 2.4 endsystem must leverage policies and mechanisms in the underlying communication infrastructure that support resource guarantees. This support can range from (1) managing the choice of the connection used for a particular invocation to (2) exploiting advanced QoS features, such as controlling the ATM virtual circuit cell pacing rate [25].

**2. OS scheduling mechanisms:** ORBs exploit OS mechanisms to schedule application-level activities end-to-end. The real-time CORBA features in CORBA 2.4 target fixed-priority real-time systems [26]. Thus, these mechanisms correspond to managing OS thread scheduling priorities. The Real-time CORBA specification in CORBA 2.4 focuses on operating systems that allow applications to specify scheduling priorities and policies. For example, the real-time extensions in IEEE POSIX 1003.1c [27] define a static priority FIFO scheduling policy that meets this requirement.

**3. Real-Time ORB endsystem:** ORBs are responsible for communicating requests between clients and servers transparently. A real-time ORB endsystem must provide standard interfaces that allow applications to specify their resource requirements to the ORB. The QoS policy framework defined by the OMG Messaging specification [19] allows applications to configure ORB endsystem resources, such as thread priorities, buffers for message queueing, transport-level connections, and network signaling, in order to control ORB behavior.

**4. Real-time services and applications:** Having a real-time ORB manage endsystem and communication resources only

provides a partial solution. Real-time CORBA ORBs must also preserve efficient, scalable, and predictable behavior end-to-end for higher-level services and application components. For example, a global scheduling service [5, 28] can be used to manage and schedule distributed resources. Such a scheduling service can interact with an ORB to provide mechanisms that support the specification and enforcement of end-to-end operation timing behavior. Application developers can then structure their programs to exploit the features exported by the real-time ORB and its associated higher-level services.

To manage the ORB endsystem capabilities outlined above, CORBA 2.4 defines standard interfaces and QoS policies that allow applications to configure and control the following resources:

- *Processor resources* via thread pools, priority mechanisms, intra-process mutexes, and a global scheduling service;
- *Communication resources* via protocol properties and explicit bindings; and
- *Memory resources* via buffering requests in queues and bounding the size of thread pools.

Applications can specify these CORBA 2.4 QoS policies along with other policies when they call standard ORB operations, such as `validate_connection` or `create_POA`. For instance, when an object reference is created using a QoS-enabled portable object adapter (POA), the POA ensures that any server-side policies that affect client-side requests are embedded within a *tagged component* in the object reference. Tagged components are name/value pairs that can be used to export attributes, such as security or QoS values, from a server to its clients within object references [2]. Clients who invoke operations on such object references implicitly use the tagged components to honor the policies required by the target object.

### 2.2.1 Evaluating Avionics Application Software Architectures

**Overview:** Figure 3 shows a conventional non-CORBA architecture for distributing periodic I/O events throughout an avionics application. This example has the following participants:

- **Aircraft sensors:** Aircraft-specific devices generate sensor data at regular intervals, *e.g.*, 30 Hz, 15 Hz, 5 Hz, *etc.* The arrival of sensor data generates interrupts that notify the mission computing applications to receive the incoming data.
- **Sensor proxies:** Mission computing systems must process data to and from many types of aircraft sensors, including global position system (GPS), inertial navigation set (INS), and forward looking infrared radar. To decouple the details

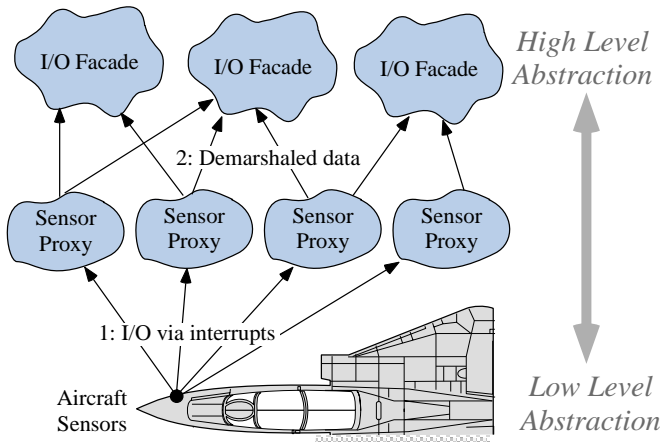


Figure 3: Example Avionics Mission Control Application

of sensor communication from the applications, sensor proxy objects are created for each sensor on the aircraft. When I/O interrupts occur, data from a sensor is given to an appropriate sensor proxy. Each sensor proxy object demarshals the incoming data and notifies I/O facade objects that depend on the sensor's data. Since modern aircraft can be equipped with hundreds of sensors, a large number of sensor proxy objects may exist in the system.

- **I/O facades:** I/O facades represent objects that depend on data from one or more sensor proxies. I/O facade objects use data from sensor proxies to provide higher-level views to other application objects. For instance, the aircraft position computed by an I/O facade is used by the navigation and weapons release subsystems.

The *push*-driven model described above is commonly used in many real-time environments [29], such as industrial process control systems and military command/control systems. One positive consequence of this push-driven model is its efficient and predictable execution of operations. For instance, I/O facades only execute when their event dependencies are satisfied, *i.e.*, when they are called by sensor proxies.

In contrast, using a *pull*-driven model to design mission applications would require I/O facades that actively acquire data from sensor proxies. If data were not available to be pulled, the calling I/O facade must block awaiting a result. Thus, for I/O facades to pull, the system must allocate additional threads to allow the application to progress while I/O facade tasks block. Adding threads to the system has many negative consequences, however, such as increased context switching overhead, synchronization complexity, and complex real-time thread scheduling policies [22]. Conversely, by using the push model, blocking is largely alleviated, which reduces the need for additional threads. Therefore, this paper focuses on the

push model.

**Drawbacks with conventional avionics architectures:** A disadvantage to the architecture shown in Figure 3 is the strong coupling between suppliers (sensor proxies) and consumers (I/O facades). For instance, to call back to I/O facades, each sensor proxy must know which I/O facades depend on its data. As a result, changes to the I/O facade layer, such as adding/removing consumers, require sensor proxy modifications. Likewise, consumers that register for callbacks are tightly coupled with suppliers. If the availability of new hardware, such as forward looking infrared radar, requires a new sensor proxy, the I/O facades must be altered to take advantage of the new technology.

**Alleviating drawbacks with the CORBA Event Service:** Figure 4 shows how the CORBA Event Service [30] can help

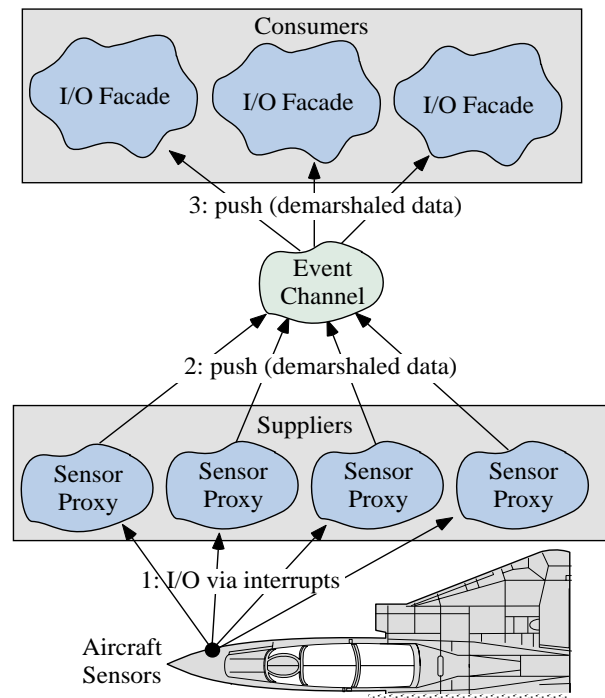


Figure 4: Example Avionics Application with CORBA Event Channel

alleviate the disadvantages of the tightly coupled consumers and suppliers shown in Figure 3. The CORBA Event Service defines *supplier*, *consumer*, and *event channel* participants so that distributed applications can exchange requests asynchronously via an *event-based* execution model [29]. Suppliers generate events, consumers process events sent by suppliers, and event channels propagate events to consumers on behalf of suppliers.

The architecture of our example avionics mission computing application [6] centers on the Publish/Subscribe [11] pattern shown in Figure 4. As shown in this figure, sensors generate an interrupt to indicate data availability. Sensor proxies then push the event to an event channel, which dispatches events to application-level consumers on behalf of the sensors generating the events. When consumers subscribe with an event channel, they indicate what types of events they are interested in receiving by supplying the filtering criteria. The benefit of using an event channel is that sensor proxies are unaffected when I/O facades are added or removed.

Before running production mission computing applications, the system is analyzed and events are assigned priorities by a real-time (RT) scheduling service [5]. Since the same consumer can receive events at different priorities, the DOC middleware must support this use-case, thereby allowing applications to use the same object reference to access a service at different priority levels.

Another benefit of a CORBA event channel-based architecture is that an I/O facade need not know which sensor proxies supply its data. Since the channel mediates on behalf of the sensor proxies, I/O facades can register for certain types of events (e.g., GPS and/or INS data arrival) without knowing which sensor proxies actually supply these types of events. Once again, the use of an event channel makes it possible to add or remove sensor proxies without changing I/O facades.

### 3 Evaluating CORBA for Real-time Applications

#### 3.1 Overview

During the past three years, we developed two versions of a real-time avionics mission computing application, each using a different version of TAO. The first version, described in [6], is based on CORBA 2.3 and uses multiple instances of the ORB, one for each priority level. This approach results in minimal priority inversion, as no resources are shared between different priority levels. However, it is hard to program to, since the servers must activate each servant multiple times, once under each ORB, and the client must manually select the right object reference for each priority level.

In the following sections we describe real-time policies and mechanisms used in our second implementation, where TAO is based on the Real-time [18] and Messaging [19] specifications. We also motivate several enhancements we developed to provide more precise control over CORBA 2.4 real-time and messaging features.

#### 3.2 Preserving Priorities End-to-End

**Context:** Systems with stringent QoS requirements, such as our avionics mission computing application, often must execute a request at the same priority, end-to-end, as described in Section ???. In the following paragraphs we outline the Real-time CORBA mechanisms in CORBA 2.4 intended to preserve request priorities end-to-end.

- **Priority mapping:** The specification defines a universal, platform-independent priority representation called the *CORBA Priority*. This feature allows applications to make prioritized CORBA invocations in a consistent fashion between nodes running on operating systems with different priority schemes. *Priority mapping functions* are used to map priority values specified in terms of *CORBA priority* into *native OS priority*.

- **CORBA priority models:** The Real-time CORBA specification defines a *PriorityModel* policy that determines the priority at which server handles requests from clients. The policy can have one of the two values: *SERVER\_DECLARED* or *CLIENT\_PROPAGATED*. In the *SERVER\_DECLARED* model shown in Figure 5 (A), the server handles requests at the pri-

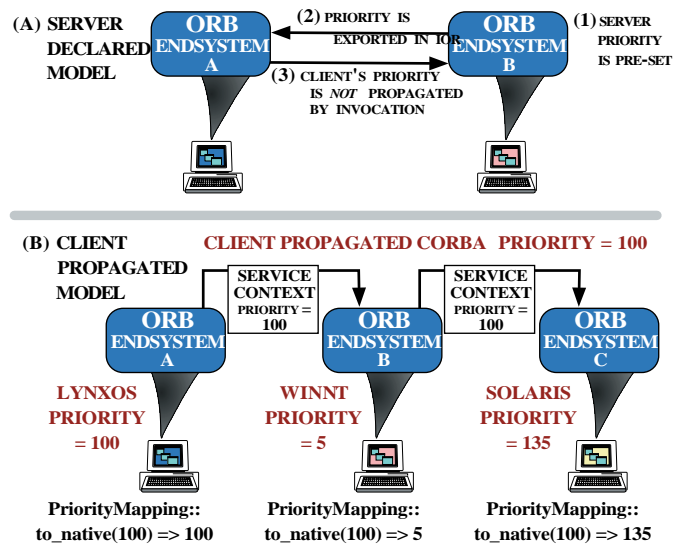


Figure 5: Real-time CORBA Priority Models

riority declared on the server side at object creation time. This priority is communicated to the client in an object reference.

The Real-time CORBA specification also defines the *CLIENT\_PROPAGATED* model shown in Figure 5 (B). In this model, the client encapsulates its priority in the service context list of the operation invocation and the server then honors the priority of the invocation. When a server ORB parses the request, it extracts the priority from the service context and sets

the priority of the processing thread to match the requested priority.

- **Thread pools:** A Real-time CORBA server can associate each POA with a pool of pre-allocated threads running at appropriate priorities. A pool can optionally be pre-configured for a maximum buffer size or number of requests, as shown in Figure 6. If buffering is enabled for the pool, the request will




Figure 6: Buffering Requests in Real-time CORBA Thread Pools

be queued until a thread is available to process it. If no queue space is available or request buffering was not specified the ORB should raise a TRANSIENT exception, which indicates a temporary resource shortage. When the client receives this exception it can reissue the request at a later point.

- **Priority banded connections:** This feature allows a client to communicate with the server via multiple transport connections. Each connection is dedicated to carrying invocations of distinct CORBA priority or range of priorities, as shown in Figure 7. A client ORB establishes a priority banded




Figure 7: Priority Banded Connections

connection by sending a server the `_bind_priority_band` request, which specifies the range of priorities the connection will be used for. This feature allows the server to allocate the necessary resources for the connection and to configure these resources to provide service for the specified priority range. The selection of the appropriate connection for each invocation is transparent to the application, and is done by the ORB based on the value of the *PriorityModel* policy.

**Problems:** As outlined above, the mechanisms defined in the Real-time CORBA chapter of the CORBA 2.4 specification provide application developers with greater control over ORB endsystem resources than earlier CORBA 2.3 specifications. For many real-time applications these mechanisms are sufficient to provide the necessary QoS guarantees. For real-time applications with stringent QoS requirements such as those outlined in Section ??, however, this lack of specificity can lead to ineffective or *non-portable* implementations, as discussed below:

- **Priority mapping problems:** Although Real-time CORBA mandates each ORB to provide default priority mapping functions, as well as a mechanism to allow users to override these defaults, it does not state how those mappings functions are accessed and set. Thus, application developers are forced to use proprietary interfaces.

- **CORBA priority model problems:** The Real-time CORBA CLIENT\_PROPAGATED model can be inappropriate for applications with hard real-time requirements due to opportunities for priority inversion [31]. In particular, it is possible that the initial priority of the thread reading the request is too high or too low, relative to the priority of the thread that processes the servant in an upcall.

Likewise, the SERVER\_DECLARED priority model is not appropriate for applications that invoke the *same* operation on the *same* object, but at *different* priorities from *different* client threads. For example, if our avionics mission computing application were to use the SERVER\_DECLARED priority model, it would have to activate the same servant multiple times, using a different priority for each activation. The client application would then choose the object reference based on the client thread's priority, and invoke the operation on the right object. However, this solution is unnecessarily complicated for the following reasons:

- It would interact poorly with CORBA location services, such as Naming or Trading, because each object must be registered multiple times.
- An application-specific client convention would be required to (1) fetch all the object references for the same object and (2) map priorities to the corresponding object references.

In our CORBA 2.3-based implementation, we faced similar challenges in managing multiple object references corresponding to multiple server ORBs. Our experiences indicated that such an approach yielded complex code that was hard to maintain, thereby negating several advantages of DOC middleware.

- **Thread pool problems:** The Real-time CORBA specification does not provide any policies to ensure that threads in a pool receive requests directly from connections. Thus,

a compliant implementation may choose to separate threads that perform all the I/O, parse the request to identify the target POA and priority, and hand off the request to the appropriate thread in the POA thread pool, as shown in Figure 8. Such an implementation can increase average and worst-case

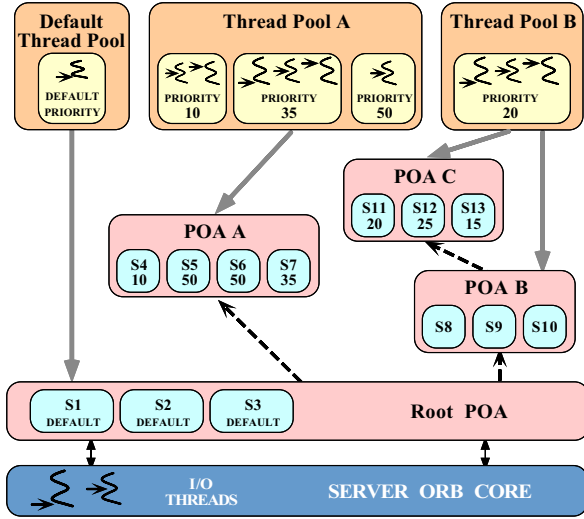


Figure 8: An Inappropriate CORBA Thread Pool Architecture for Hard Real-time Applications

latency and create opportunities for unbounded priority inversions [15], however. For instance, even under a light load, the server ORB incurs a dynamic memory allocation, multiple synchronization operations, and a context switch to pass a request between a network I/O thread and a POA thread.

- **Priority banded connection problems:** There is no standard API in Real-time CORBA that allows server applications to control how thread pools are associated with priority banded connections. For instance, a server application can not control whether its ORB assigns each connection a separate thread, or whether a pool of threads can be pre-allocated to service multiple connections that have the same priority range. Unfortunately, this lack of detail in the specification makes it hard to write real-time applications that behave predictably across different ORB platforms.

The Real-time CORBA specification also lacks a standard API that would allow a server application to control how its ORB associates a thread at a pre-specified priority to read requests from a priority banded connection. Thus, the actual ORB thread that performs I/O operations could be different from the thread processing the request, and could execute at the wrong priority, thereby incurring priority inversion. This lack of specificity in the Real-time CORBA priority banded connections mechanism can lead to implementations that suffer from problems similar to those with POA thread pools shown in Figure 8.

**Solution → Prioritized connection endpoints:** To alleviate the problems listed above, we defined mechanisms in TAO to explicitly map thread pools and thread priorities to connection endpoints. These mechanisms extend the Real-time CORBA specification to give TAO applications greater control over the mapping of connections to thread priorities within the ORB Core. Figure 9 provides an example of the avionics mission computing application shown in Figure ?? revised to use TAO’s prioritized connection endpoints mechanism.

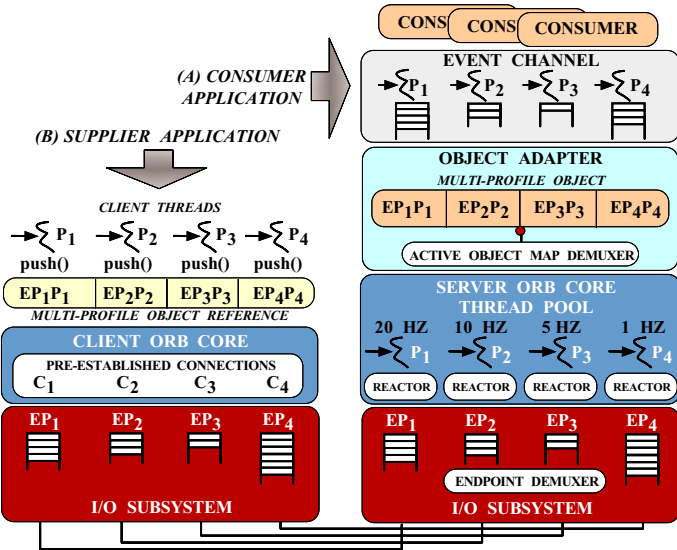


Figure 9: TAO’s New CORBA 2.4-based Solution

As shown in this figure, the server application is an event channel consumer with four connection endpoints,  $EP_1 \dots EP_4$ . Each endpoint is assigned a CORBA priority, e.g.,  $EP_1$  has priority  $P_1$ , and is serviced by a thread of the corresponding native OS priority. Object references for servants activated in this server contain four profiles, one for each endpoint, as shown in the object adapter portion of the server in Figure 9 (A).

The client application in Figure 9 (B) is an event channel supplier. It has four threads with priorities  $P_1 \dots P_4$ . When a supplier thread makes a call on the consumer object reference exported by the server, the ORB finds a pre-established connection for that endpoint and uses it to send the request. The priority is preserved end-to-end because on the server-side the connection is serviced by a thread at the same priority as the thread making the request on the client-side.

**Implementing prioritized connection endpoints in TAO:** Below, we describe key mechanisms provided by TAO to implement prioritized connection endpoints in client and server ORBs:

- **Server ORB support for binding thread priorities to listen-mode connection endpoints:** TAO allows servers to



have multiple listen-mode [32] connection endpoints, each associated with a CORBA priority. Each connection endpoint is also statically associated with a pool of threads running at native OS priority corresponding to the CORBA priority of the endpoint. Pool threads are responsible for accepting and servicing connections on the associated endpoint.

When an object is activated in a server with multiple endpoints, the generated object reference contains multiple profiles, one for each endpoint. Each profile stores the CORBA priority of its endpoint in a tagged component. This design allows a client to receive service at its desired priority by simply selecting and using the profile containing that priority.

TAO's prioritized connection endpoints extension is particularly attractive for applications, such as avionics mission computing, that invoke operations on the same object at different priority levels. In particular, these applications can specify a set of prioritized connection endpoints on the command-line, effectively defining the set of priorities supported on the server *a priori*, thereby allowing the ORB to schedule and allocate resources more effectively end-to-end. Moreover, this programming model is much simpler than creating multiple objects and object references and trying to assign them different thread priorities.

- **Client ORB support for connections with priorities:**

When a client makes an invocation, the client-side ORB must select one of the profiles from an object reference before sending the request to the server. The profile is selected based on the priority at which the client wants the request serviced. To allow clients to specify this desired priority, TAO defines a *ClientPriority* policy. Clients can set the *ClientPriority* policy to one of the following values:

- `USE_NO_PRIORITY` – *i.e.*, priority information is not used when a client ORB selects a profile from an object reference.
- `USE_THREAD_PRIORITY` – *i.e.*, the priority of the client thread sending a request is used to select the profile. This option is used when the priority of request must be preserved end-to-end. For example, we use this option in Figure 9 (B), where the *ClientPriority* policy is checked before the client ORB selects a profile from the object reference. In that case, when a client thread with priority  $P_1$  invokes an operation on a consumer object, the ORB selects the profile corresponding to that priority, *i.e.* the one that contains connection endpoint  $EP_1$ .
- `USE_PRIORITY_RANGE` – In this case, a range of priorities to be used for profile selection is specified by the application inside the policy. This option allows applications to request services at a priority that is *different* than that of the client thread invoking an operation. For example, a high-priority client thread can generate a

low-priority event, such as a display update. This high-priority thread can post the message in a remote server at a low-priority, to minimize the effect on more critical processing. However, it need not change its own priority to perform this task, which avoids local priority inversions.

TAO's prioritized connection endpoints and *ClientPriority* policy extend the standard Real-time CORBA priority models and its priority banded connections mechanism to achieve an effective balance between the `SERVER_DECLARED` and `CLIENT_PROPAGATED` models. In particular, TAO provides the same degree of control to the server as the `SERVER_DECLARED` model by restricting clients to use well-known priorities. However, it also allows clients to select a priority published by the server that best meets their requirements. TAO's design avoids priority inversions and ensures ORB endsystem resources are strictly controlled, while still retaining a simple programming model.

Section 4.1 illustrates the performance of TAO's prioritized connection endpoint architecture.

### 3.3 Achieving Reliable Asynchronous Communication

**Context:** Embedded real-time CORBA applications often use one-way operations to simulate message-passing via standard CORBA features. For example, avionic mission computing applications [6] process periodic event messages, such as sensor updates and heartbeat messages from redundant systems. Typically, clients send these messages to servers via CORBA one-way operations, which require no response.

**Problems:** The semantics of conventional CORBA one-way operations are often unacceptable because the CORBA 2.3 specification does not require an ORB to guarantee that one-way operations will be delivered [4].

**Solution → CORBA 2.4 reliable asynchronous features:** To alleviate the problem outlined above, the CORBA Messaging specification defines a policy called *SyncScope* that allows clients more control over the degree of reliability for one-way operation invocations. Figure 10 illustrates the following four levels of reliability for one-way operations:

- `SYNC_NONE`: With this policy value, the client ORB returns control to the client application before passing the request to the transport layer. This value minimizes the amount of time a client spends blocking on the one-way operation, but provides the lowest level of delivery guarantee. The `SYNC_NONE` policy is useful for applications that require minimal client operation latency, while tolerating reduced reliability guarantees.

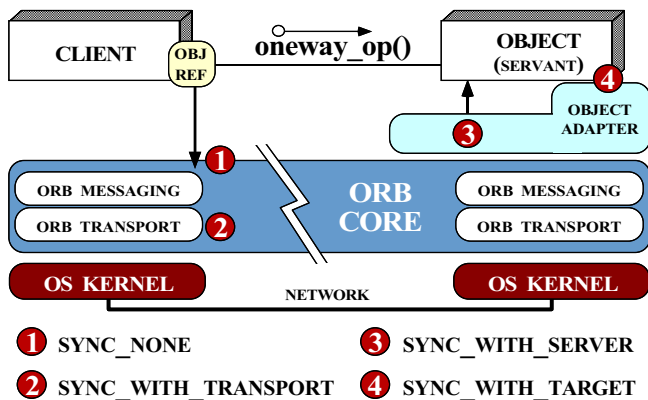


Figure 10: Reliable One-way Synchronization Scopes

- **SYNC\_WITH\_TRANSPORT:** With this policy value, the ORB returns control to the client only after the request is passed successfully to the transport layer, *e.g.*, the client’s TCP protocol stack. A client can incur unbounded latencies if a connection endpoint flow controls due to a limited buffer space. When used with a connection-oriented transport, such as TCP, SYNC\_WITH\_TRANSPORT can provide more assurance than SYNC\_NONE. This policy is appropriate for clients that require a compromise between low latency and reliable delivery.

- **SYNC\_WITH\_SERVER:** With this policy value, the client invokes a one-way operation and then blocks until the server ORB sends an acknowledgment. The server ORB sends the acknowledgment after invoking any servant managers, but before dispatching the request to the servant. The SYNC\_WITH\_SERVER policy value provides clients with assurance that the remote servant has been located. This feature is particularly useful for real-time applications that require some degree of reliability, *e.g.*, because they run over backplanes that lose packets occasionally, but need not wait for the entire servant upcall to complete.

- **SYNC\_WITH\_TARGET:** This policy value is equivalent to a synchronous two-way CORBA operation, *i.e.*, the client will block until the server ORB sends a reply after the target object has processed the operation. If no exceptions are raised, the client can assume that the target servant processed its request. This synchronization level is appropriate for clients that need assurance that the upcall was performed and can tolerate the additional latency.

**Implementing reliable one-ways in TAO:** The *SyncScope* policy controls the reliability of one-way requests. It can be set at the object-level, thread-level, or ORB-level. As with any CORBA policy, the more specific levels can override the more general levels. To implement reliable one-way requests, TAO’s IDL compiler [33] generates client stub code

that checks the *SyncScope* policy value and sets the appropriate bits in the *response\_flags* field in the GIOP request header.

If the *SyncScope* policy is SYNC\_NONE, the request is buffered, as described in Section 3.4. If the policy value is SYNC\_WITH\_SERVER or SYNC\_WITH\_TARGET, the client ORB must wait for a reply from the server and check for a LOCATION\_FORWARD response or a CORBA system exception.

On the server, the ORB’s behavior is based solely on the value of the *response\_flags* field of the request header. If the flags are set to a *SyncScope* policy value of SYNC\_WITH\_TARGET, the request is treated as a two-way request, whether it originated as a one-way or as a two-way. If the flags are set to a value of SYNC\_WITH\_SERVER, however, a response will be initiated by the Object Adapter immediately after it locates the servant, but before dispatching the upcall.

Section ?? presents benchmarks illustrating the performance of TAO’s reliable one-way implementation.

### 3.4 Ensuring Adequate Operation Throughput

**Context:** Distributed real-time applications often have stringent timing requirements, where critical operations must begin and/or complete within specified time intervals. For example, aircraft sensor devices, such as navigation devices and radar sensors, generate data that must be processed at regular periodic intervals [6]. Such applications often have a fixed time period in which to invoke remote one-way operations. After invoking each operation, the client must perform other processing, *i.e.*, it does not wait synchronously for the server to process the operation and respond.

**Problems:** The following two problems can arise when applying ORB middleware to distributed real-time applications with periodic processing requirements:

- **Inadequate operation throughput:** The time spent delivering a one-way or asynchronous operation to a server includes the overhead of invoking one or more *write* calls to the client OS. In turn, this incurs protocol stack and network interface processing, as well as the propagation delay across the communication media. This per-operation overhead constitutes a non-trivial amount of the total end-to-end latency, particularly for small requests. As a result of this overhead, real-time applications may be limited to a relatively low number of remote operations per time period.

- **Blocking flow control:** It is important that periodic real-time applications not block indefinitely when ORB endsystem and network resources are unavailable temporarily. However, ORB transport protocols, such as IIOP, often implement reliable data delivery using a sliding window flow control algorithm [32]. Thus, they may block the client from transmitting

additional data when the communication channel is congested, or if the server is slow. Although some transport protocols buffer a limited number of bytes or requests, they will typically block client threads after this limit is reached.

One way to solve these problems is to revise the application’s IDL interfaces and reimplement the clients so they buffer data at the application-level. This solution works well for certain periodic applications that can sacrifice some latency for increased operation throughput. Buffering at the application-level increases the burden on application developers, however, thereby increasing the implementation, validation, and maintenance effort. Moreover, if two or more application IDL interfaces require buffering, code can be duplicated unnecessarily, which increases application footprint.

**Solution • ORB-level request buffering:** Often, a more effective solution is to have the ORB buffer one-way and asynchronous invocations transparently.<sup>3</sup> At some later time, the buffered requests can be delivered *en masse* to the server. There are several benefits to ORB-level request buffering:

- By buffering requests, a client ORB amortizes the per-operation processing overhead and increases effective network utilization.
- The ORB can use OS *gather write* operations, such as `writew` [32], to minimize the number of mode switches needed to transmit the buffered requests.
- ORB-level buffering can increase application control over the buffering of CORBA requests. This feature is important when the buffering provided by the transport protocol is inadequate, thereby forcing indefinite blocking of the client due to flow control.

**Implementing ORB-level request buffering in TAO:** The CORBA 2.4 Messaging specification introduces several mechanisms to give application developers more control over QoS parameters than in CORBA 2.3 specifications. In particular, applications can use CORBA 2.4 features, such as the *Sync-Scope* policy, to control latency/reliability tradeoffs. For example, applications can use the `SYNC_WITH_SERVER` policy value to achieve reliable transport delivery, without waiting for the entire servant’s computation to complete. Likewise, the application can ensure non-blocking behavior by using the `SYNC_NONE` policy value, which TAO implements by buffering multiple requests before sending them to the server.

Unfortunately, neither the CORBA Messaging or the RT CORBA specifications provide mechanisms to control the size or duration of buffers, nor does it provide explicit interfaces to

<sup>3</sup>Synchronous two-way requests and reliable one-way operations should not be buffered. The ORB must deliver these request immediately to the server because the client waits for the server’s response before continuing.

flush buffers. These semantics are insufficient for applications that require precise control over the ORB utilization of memory and network resources. Therefore, we have extended TAO to allow applications to specify multiple strategies for delivering buffered requests via a new *BufferingConstraint* policy.

Figure 11 illustrates how TAO uses this policy to buffer one-way invocations inside the ORB Core for subsequent delivery. When application-specified buffering limits are reached, the

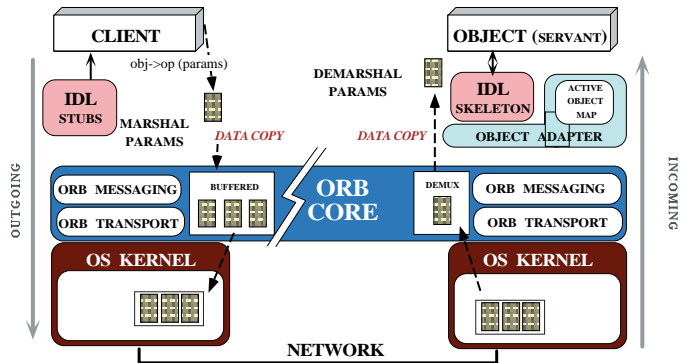


Figure 11: One-way and Asynchronous Request Buffering

buffers are flushed and the queued requests are delivered to the server.

A combination of the following conditions can be specified simultaneously using TAO’s *BufferingConstraint* policy:

- 1. Message Count:** When the number of buffered messages reaches an application-specified high-water mark, the buffered requests are delivered to the server. This approach allows applications to batch  $n$  requests together.
- 2. Message Bytes:** When the number of bytes in the buffered messages reaches an application-specified high-water mark, the buffered requests are delivered to the server. This approach allows applications to buffer  $n$  bytes at the ORB layer.
- 3. Periodic Timeout:** After an application-specified time interval, the ORB delivers any buffered requests to the server. This approach allows applications to pace the delivery of messages to the server even when the requests are produced at irregular intervals.
- 4. Explicit Flushing:** Applications can flush any queued messages explicitly. This approach allows applications to deliver the batched messages to the server in response to some external event.
- 5. Out-of-Band Requests:** Applications can skip buffering for some requests. This approach allows applications to deliver urgent requests to the server immediately, bypassing the buffered requests.

Section 4.2 presents benchmarks illustrating the performance of TAO’s buffered requests implementation.

### 3.5 Evaluation of the CORBA 2.4-based Solution

Our CORBA 2.4-based avionics mission computing solution has the following improvements over the original CORBA 2.3-based design:

**More standard programming model:** The CORBA 2.4 Real-time specification defines a standard model for implementing many features required for avionics mission computing using only a single ORB per CPU. This model supports the mapping of priorities to particular invocations, objects, and threads. As a result, application programming is simplified and the portability of application software increases because the system is based on a standard.

**Reduced memory footprint:** Our original CORBA 2.3-based solution required multiple ORBs be created within each client and server, once for each rate group. Moreover, servants had multiple object references, one for each global priority. This design resulted in a relatively large application footprint. In contrast, our CORBA 2.4-based solution allows each client and server application to create a single ORB. Thus, only one object reference is created per servant, which further reduces the overall memory footprint.

**Efficient initialization:** In addition to simplifying the programming model and minimizing the required memory resources, the use of one ORB per-process reduces the time required to initialize the avionics mission computing applications. Reducing this overhead is particularly important when the system must recover from transient power cycles.

**Simplified client threading model:** Clients are greatly simplified because they manage only one set of object references. Policies and object references contain sufficient information for the ORB to determine the appropriate connection to use on each request.

**Improved priority preservation:** By supporting multiple connection endpoints within server ORBs, the CORBA 2.4-based implementation has several benefits. For example, the destination service access point, such as the TCP port number, can be mapped to a global CORBA priority thereby ensuring that all CORBA requests within a connection queues have the same priority. This early demultiplexing [34] technique, combined with client and server ORBs' respect of a request's priority, results in *vertically* (*i.e.*, network interface ↔ application layer) and *horizontally* (*i.e.*, peer-to-peer) integrated ORB endsystems. The resulting DOC middleware environment preserves invocation priorities end-to-end, *i.e.*, throughout the ORB endsystems and inter-ORB connections.

**Improved one-way invocation semantics:** Real-time applications must often balance the competing needs of reliable communications, network throughput, and invocation latency.

The improved semantics added by the CORBA 2.4 Messaging specification [19] gives clients greater control over these tradeoffs.

In addition to alleviating the drawbacks with our original CORBA 2.3-based solution, the new CORBA 2.4-based version of TAO also provides the following benefits:

**Easier integration with CORBA common object services:** To use CORBA common object services, such as Naming and Trading, in the CORBA 2.3-based approach, a server must export multiple object references to the same servant, one for each priority. Then, to locate the object reference corresponding to the desired priority, a client must use an *ad hoc* mechanism to retrieve the desired object references at the appropriate priorities. In contrast, in the CORBA 2.4-based approach there is no need for multiple object references. Therefore, no *ad hoc* protocol for mapping priorities to object references is required.

With other services, CORBA 2.3-based approach may require modification to the service itself. For example, the Event Service invokes operations on application-provided objects. To invoke these operations at the appropriate priorities the service must: (1) have access to multiple object references for each application object and (2) select object references corresponding to desired priorities using the application defined protocol. In contrast, in the CORBA 2.4-based approach these tasks are all performed by the ORB transparently to applications and services.<sup>4</sup>

**Easier integration with real-time scheduling services:** The CORBA 2.4 Real-time specification supports higher-level CORBA scheduling services that allocate resources end-to-end. For example, TAO's static scheduling service [5] can associate application activities with global CORBA priorities. Such scheduling services can use the priority transformations and the policy framework defined in the CORBA Messaging [19] to create sophisticated and adaptive real-time applications.

## 4 The Performance of the TAO Real-time CORBA ORB

### 4.1 Preserving Priorities End-to-End

**Overview:** The benchmarks in this section compare the performance of TAO's CORBA 2.3-based solution, *i.e.*, using an ORB-per-priority architecture, with TAO's CORBA 2.4-based solution, *i.e.*, using the prioritized connection mechanism described in Section 3.2. In particular, to determine how well

<sup>4</sup>It is still possible that implementations of these services are unsuitable for real-time application, *e.g.*, due to excessive priority inversion.

each approach preserves priorities end-to-end, we compare the latency and jitter of a high-priority client thread as it competes with a variable number of low-priority client threads.

**Hardware/OS Benchmarking Platforms:** All benchmarks in this section were performed between two 266 MHz PowerPC boards with 32 MBytes of RAM, running the LynxOS 3.0.0 operating system and connected by a 100 Mbps Ethernet. The tests were run with real-time, preemptive, FIFO thread scheduling, which provides strict priority-based scheduling to application threads.

**Measurement techniques:** Below we describe the client-side and the server-side parts of the benchmark.

- **Client-side:** On the client, a single high-priority thread and a variable number of low-priority threads run concurrently. Both CORBA 2.3-based solution and CORBA 2.4-based solution were benchmarked with 1, 3, 6, 9, 12, and 15 low-priority client threads. Each low-priority thread has a different priority value. The range of LynxOS native priorities used by these threads is 64 to 79. The high-priority thread runs at priority 128.

When the test program creates the client threads, these threads block on a barrier lock so that no client thread starts until the others are created and are ready to run. When all client threads are ready to send requests, the main thread unblocks them. Each client thread issues 20,000 requests to the server at the fastest possible rate.

- **Server-side:** On the server, a servant is created and configured to service client requests at the same priorities as those of its client peers. In the original CORBA 2.3-based approach, this is achieved by creating an ORB-per-thread for each test priority. In the CORBA 2.4-based approach, this is achieved by mapping connection endpoints to threads possessing the appropriate priorities.

Regardless of the approach, each request is serviced at the same priority as that of the invoking client thread. This mapping is achieved using different mechanisms in the two configurations. In CORBA 2.3 configuration, each client thread uses a different object reference, *i.e.*, the one published by the server ORB running in the thread with corresponding priority. In CORBA 2.4 configuration, the client sets the value of the *ClientPriority* policy to `USE_THREAD_PRIORITY`.

As was mentioned earlier, each low-priority thread has a different priority value, *e.g.*, 64, 65, etc. Such system configuration is more demanding than simply having all low-priority threads run at the same priority. By using different priority values for the low-priority threads, concurrency can be increased on the server. Thus, this design provides more opportunities for priority inversion to occur if the underlying ORB is not designed and implemented properly.

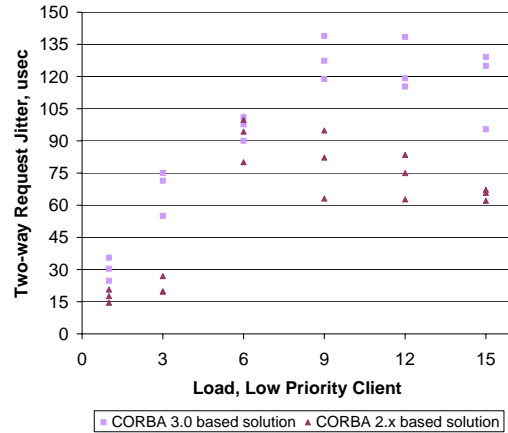


Figure 12: Jitter for High-Priority Client Thread

**Results:** Each solution was benchmarked with a different number of low-priority client threads: 1, 3, 6, 9, 12, and 15. For each solution and each number of client threads, the experiment was repeated three times. Figure 12 shows the jitter results. Figure 13 shows average latency (over three samples), with average jitter shown as error bars. The results in these




Figure 13: Latency for High-Priority Client Thread

figures illustrate that the CORBA 2.4-based latency is slightly lower than the CORBA 2.3 version, though its jitter is slightly higher. In general, therefore, the CORBA 2.4-based configuration provides a simpler programming model, smaller footprint, and faster initialization, with negligible impact on efficiency and predictability.

## 4.2 Buffered Request Benchmarks

**Overview:** We conducted benchmarks to compare the throughput of various buffered one-way request configurations. The *SyncScope* policy value was set to `SYNC_NONE` and the maximum message count of the buffer was increased for each test run.

**Hardware/OS Benchmarking Platforms:** We used two Dell computers, each with four Pentium Xeon CPUs running at 400 MHz and connected by a 100 bps Ethernet. The operating system of both machines was Windows NT 4.0. This OS, unlike most UNIX systems, does not limit the value of `IOV_MAX`, which is the maximum size of the vector used in *gather-write* operations. Thus, Windows NT is better equipped to demonstrate the full potential of performance gains from buffering one-way requests in the ORB. The timer used on both platforms was the high-resolution timer of the Intel chip, accessed by the Pentium RDTSC instruction.

**Measurement techniques:** A timestamp was obtained immediately before and after the execution of the request iteration loop. Throughput values were then calculated in calls/second for each test run of 2,000 requests. To closely emulate a real-time transport mechanism, such as VME, certain configuration policies were followed when conducting the one-way request benchmarks. In particular, Nagle’s algorithm was disabled and the TCP window size was set to 8 kilobytes, which was the minimum supported on Windows NT. The size of each request was set to a larger value (9 kilobytes plus the header size) so that TCP would flow-control on each request and not buffer additional requests.

**Results:** In Figure 14, the throughput of unbuffered one-way requests is shown in the leftmost bar and is used as a baseline for comparison with the throughput of buffered one-way requests. To measure the effect of the buffer size, we set it to

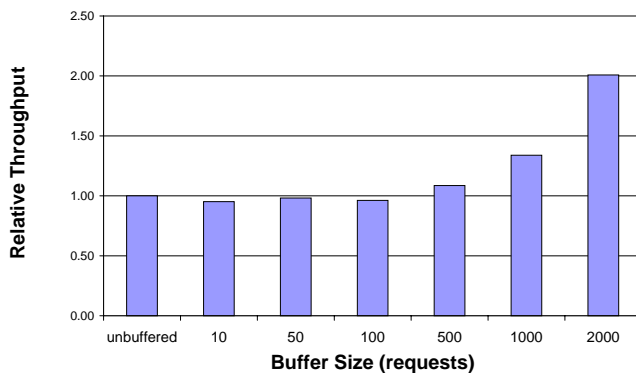


Figure 14: Throughput for Buffered One-way Requests

a different value in each test run. The other variables remain

unchanged, *i.e.*, in all test runs the same number of messages is sent, and all the messages are identical. The figure shows that significant throughput gains (over 100% in the rightmost bar) can be realized by TAO’s request buffering mechanism. However, tuning the queue size for each use-case is crucial to maximize performance.

## 5 Concluding Remarks

Real-time distributed object computing (DOC) middleware is a promising solution for key challenges facing researchers and developers of real-time applications. Designing and optimizing standards-based and commodity-off-the-shelf (COTS) DOC middleware that can meet the QoS requirements of real-time applications requires an integrated architecture that can deliver end-to-end QoS support at multiple levels in real-time and embedded systems. The Real-time CORBA [18] and Messaging [19] specifications in CORBA 2.4 are an important step in this direction.

Unfortunately, the CORBA 2.4 specification lacks sufficient specificity to *portably* manage processor, communication, and memory resources for applications with stringent QoS requirements. For example, Section 3.2 describes outlines a number of problems with the standard Real-time CORBA *priority mapping*, *priority model*, *thread pool*, and *priority banded connection* mechanisms intended to preserve request priorities end-to-end. The following discussion summarizes our lessons learned developing and deploying an implementation of CORBA and representative applications that are based on the CORBA 2.4 specifications.

**Control over thread priority must be end-to-end:** Ensuring appropriate end-to-end QoS for real-time applications requires more than just the implementation of the Real-time CORBA specification and the mapping of global priorities between ORBs. It also requires control over end-to-end thread priorities, connection  $\leftrightarrow$  thread pool associations, and other ORB endsystem resources used to process a request. Moreover, applications must be able to control these resources at each layer of an DOC middleware. The Real-time CORBA specification [18] adopted for CORBA 2.4 allows implementations to provide end-to-end guarantees, but it does not require them. Moreover, it does not provide explicit mechanisms to control how I/O threads in an ORB Core map to thread priorities and connections.

Section 3.2 describes the policies and mechanisms we used in TAO’s prioritized connection endpoints extension to precisely define these associations and ORB configurations. Application developers can use this extension to ensure that the client and server ORBs process the request at the appropriate priority end-to-end. In previous work [34], we described how TAO can be integrated with an *early demultiplexing* feature in

the operating system's I/O subsystem to ensure that the OS kernel and network interfaces also preserve the priority of the request.

**One-way semantics must be specified precisely:** The CORBA 2.3 specification defines the semantics of one-way operations to receive best-effort service. However, it provides no further end-to-end delivery guarantees to applications. Thus, a CORBA 2.3-compliant ORB can simply drop every request, deliver it at some arbitrary future time (but out of order with subsequent requests), or send it immediately. The CORBA 2.4 Messaging specification provides better control over the semantics of oneway operations, but it falls short in the SYNC\_NONE case, where it does not provide mechanisms to control buffering limits and/or flushing policies. Section 3.4 describe the policies and mechanisms we used in TAO to define one-way buffering semantics precisely and ensure adequate throughput for certain types of real-time applications.

The open-source code, benchmarks, and documentation for TAO is freely available and can be downloaded from URL [www.cs.wustl.edu/~schmidt/TAO.html](http://www.cs.wustl.edu/~schmidt/TAO.html). Our focus on the TAO project has been to research, develop, and optimize policies and mechanisms that allow CORBA to support applications with hard real-time requirements. These requirements motivate many of the optimizations and design strategies presented in this paper.

TAO has been used on a wide range of distributed real-time and embedded systems, including an avionics mission computing architecture for Boeing [6], the next-generation Run Time Infrastructure (RTI) implementation for the Defense Modeling and Simulation Organization's (DMSO) High Level Architecture (HLA) [35], and high-energy physics experiments at SLAC [36] and CERN [37].

## References

- [1] A. Gokhale and D. C. Schmidt, "Optimizing a CORBA IIOP Protocol Engine for Minimal Footprint Multimedia Systems," *Journal on Selected Areas in Communications special issue on Service Enabling Platforms for Networked Multimedia Systems*, vol. 17, Sept. 1999.
- [2] Object Management Group, *The Common Object Request Broker: Architecture and Specification*, 2.4 ed., Oct. 2000.
- [3] Sun Microsystems, Inc, *Java Remote Method Invocation Specification (RMI)*, Oct. 1998.
- [4] M. Henning and S. Vinoski, *Advanced CORBA Programming With C++*. Reading, Massachusetts: Addison-Wesley, 1999.
- [5] D. C. Schmidt, D. L. Levine, and S. Mungee, "The Design and Performance of Real-Time Object Request Brokers," *Computer Communications*, vol. 21, pp. 294–324, Apr. 1998.
- [6] T. H. Harrison, D. L. Levine, and D. C. Schmidt, "The Design and Performance of a Real-time CORBA Event Service," in *Proceedings of OOPSLA '97*, (Atlanta, GA), pp. 184–199, ACM, October 1997.
- [7] J. A. Zinky, D. E. Bakken, and R. Schantz, "Architectural Support for Quality of Service for CORBA Objects," *Theory and Practice of Object Systems*, vol. 3, no. 1, pp. 1–20, 1997.
- [8] I. Pyarali, T. H. Harrison, and D. C. Schmidt, "Design and Performance of an Object-Oriented Framework for High-Performance Electronic Medical Imaging," *USENIX Computing Systems*, vol. 9, November/December 1996.
- [9] C. O'Ryan, D. C. Schmidt, D. Levine, and R. Noseworthy, "Applying a Scalable CORBA Events Service to Large-scale Distributed Interactive Simulations," in *Proceedings of the 5<sup>th</sup> Workshop on Object-oriented Real-time Dependable Systems*, (Monterey, CA), IEEE, Nov. 1999.
- [10] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*. Reading, Massachusetts: Addison-Wesley, 1995.
- [11] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, and M. Stal, *Pattern-Oriented Software Architecture – A System of Patterns*. New York: Wiley and Sons, 1996.
- [12] DARPA, "The Quorum Program." [www.darpa.mil/ito/research/quorum/index.html](http://www.darpa.mil/ito/research/quorum/index.html), 1999.
- [13] K. Kim and E. Shokri, "Two CORBA Services Enabling TMO Network Programming," in *Fourth International Workshop on Object-Oriented, Real-Time Dependable Systems*, IEEE, January 1999.
- [14] D. C. Schmidt, M. Stal, H. Rohnert, and F. Buschmann, *Pattern-Oriented Software Architecture: Patterns for Concurrent and Networked Objects, Volume 2*. New York: Wiley & Sons, 2000.
- [15] I. Pyarali, C. O'Ryan, D. C. Schmidt, N. Wang, V. Kachroo, and A. Gokhale, "Using Principle Patterns to Optimize Real-time ORBs," *IEEE Concurrency Magazine*, vol. 8, no. 1, 2000.
- [16] M. Fayad, R. Johnson, and D. C. Schmidt, eds., *Implementing Application Frameworks: Object-Oriented Frameworks at Work*. New York: Wiley & Sons, 1999.
- [17] D. C. Schmidt, "ACE: an Object-Oriented Framework for Developing Distributed Applications," in *Proceedings of the 6<sup>th</sup> USENIX C++ Technical Conference*, (Cambridge, Massachusetts), USENIX Association, April 1994.
- [18] Object Management Group, *Real-time CORBA Joint Revised Submission*, OMG Document orbos/99-02-12 ed., March 1999.
- [19] Object Management Group, *CORBA Messaging Specification*. Object Management Group, OMG Document orbos/98-05-05 ed., May 1998.
- [20] R. Callison, M. Goo, and D. Butler, "Real-time CORBA Trade Study," Tech. Rep. D204-31159, The Boeing Company, 1999.
- [21] A. B. Arulanthu, C. O'Ryan, D. C. Schmidt, M. Kircher, and J. Parsons, "The Design and Performance of a Scalable ORB Architecture for CORBA Asynchronous Messaging," in *Proceedings of the Middleware 2000 Conference, ACM/IFIP*, Apr. 2000.
- [22] 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," *Journal of Real-time Systems, special issue on Real-time Computing in the Age of the Web and the Internet*, vol. 21, no. 2, 2001.
- [23] Object Management Group, *The Common Object Request Broker: Architecture and Specification*, 2.3 ed., June 1999.
- [24] E. Eide, K. Frei, B. Ford, J. Lepreau, and G. Lindstrom, "Flick: A Flexible, Optimizing IDL Compiler," in *Proceedings of ACM SIGPLAN '97 Conference on Programming Language Design and Implementation (PLDI)*, (Las Vegas, NV), ACM, June 1997.
- [25] Z. D. Dittia, G. M. Parulkar, and J. R. Cox, Jr., "The APIC Approach to High Performance Network Interface Design: Protected DMA and Other Techniques," in *Proceedings of INFOCOM '97*, (Kobe, Japan), pp. 179–187, IEEE, April 1997.
- [26] D. C. Schmidt and F. Kuhns, "An Overview of the Real-time CORBA Specification," *IEEE Computer Magazine, Special Issue on Object-oriented Real-time Computing*, vol. 33, June 2000.
- [27] "Information Technology – Portable Operating System Interface (POSIX) – Part 1: System Application: Program Interface (API) [C Language]," 1995.
- [28] C. D. Gill, D. L. Levine, and D. C. Schmidt, "The Design and Performance of a Real-Time CORBA Scheduling Service," *Real-Time Systems, The International Journal of Time-Critical Computing Systems, special issue on Real-Time Middleware*, vol. 20, March 2001.

- [29] R. Rajkumar, M. Gagliardi, and L. Sha, "The Real-Time Publisher/Subscriber Inter-Process Communication Model for Distributed Real-Time Systems: Design and Implementation," in *First IEEE Real-Time Technology and Applications Symposium*, May 1995.
- [30] Object Management Group, *CORBAServices: Common Object Services Specification, Revised Edition*, 95-3-31 ed., Mar. 1995.
- [31] M. de Sousa, "Mapping Synchronisation Protocols onto Real-Time CORBA," in *Proceedings of the 2<sup>nd</sup> International Symposium on Distributed Objects and Applications (DOA 2000)*, (Antwerp, Belgium), OMG, Sept. 2000.
- [32] W. R. Stevens, *UNIX Network Programming, Volume 1: Networking APIs: Sockets and XTI, Second Edition*. Englewood Cliffs, NJ: Prentice-Hall, 1998.
- [33] A. B. Arulanthu, C. O’Ryan, D. C. Schmidt, and M. Kircher, "Applying C++, Patterns, and Components to Develop an IDL Compiler for CORBA AMI Callbacks," *C++ Report*, vol. 12, Mar. 2000.
- [34] F. Kuhns, D. C. Schmidt, C. O’Ryan, and D. Levine, "Supporting High-performance I/O in QoS-enabled ORB Middleware," *Cluster Computing: the Journal on Networks, Software, and Applications*, vol. 3, no. 3, 2000.
- [35] F. Kuhl, R. Weatherly, and J. Dahmann, *Creating Computer Simulation Systems*. Upper Saddle River, New Jersey: Prentice Hall PTR, 1999.
- [36] SLAC, "BaBar Collaboration Home Page." [www.slac.stanford.edu/BFROOT/](http://www.slac.stanford.edu/BFROOT/).
- [37] A. Kruse, "CMS Online Event Filtering," in *Computing in High-energy Physics (CHEP 97)*, (Berlin, Germany), Apr. 1997.