Designing and Optimizing a Scalable CORBA Notification Service

Pradeep Gore and Ron Cytron

{pradeep, cytron}@cs.wustl.edu
Department of Computer Science
Washington University, St.Louis

Douglas Schmidt and Carlos O'Ryan

{schmidt, coryan}@uci.edu
Electrical & Computer Engineering
University of California, Irvine*

Abstract

Many distributed applications require a scalable event-driven communication model that decouples suppliers from consumers and simultaneously supports advanced quality of service (QoS) properties and event filtering mechanisms. The CORBA Notification Service provides a publish/subscribe mechanism that is designed to support scalable event-driven communication by routing events efficiently between many suppliers and consumers, enforcing various QoS properties (such as reliability, priority, ordering, and timeliness), and filtering events at multiple points in a distributed system.

This paper provides several contributions to research on scalable notification services. First, we present the CORBA Notification Service architecture and illustrate how it addresses limitations with the earlier CORBA Event Service. Second, we explain how we addressed key design challenges faced when implementing the Notification Service in TAO, which is our high-performance, real-time ORB. We discuss the optimizations used to improve the scalability of TAO's Notification Service. Finally, we present empirical results of the performance of our implementation.

1 Introduction

Many distributed applications, such as real-time avionics mission computing systems, distributed interactive simulations, and computer-assisted stock trading, require an event-based communication model. The CORBA [1] Notification Service provides developers of these applications with a standards-based, QoS- and filtering-enabled event distribution mechanism. Our work on the Notification Service leverages the experience we gained developing TAO's Real-Time Event Service [2] to provide a flexible, extensible, and predictable implementation. In this paper, we explore the key design challenges faced in providing a *scalable* notification service. We also discuss optimization issues related to footprint reduction and configurability.

The CORBA Event Service: The CORBA Event Service provides provides decoupled, asynchronous and transparent communication between participants.It defines three roles:

- *Suppliers*, which produce event data;
- *Consumers*, which receive and process event data;
- *Event channels*, which are mediators [3] through which multiple consumers and suppliers communicate asynchronously.

However, the CORBA Event Service has limitations - It does not provide interfaces or policies for supporting QoS properties such as reliability, priority, ordering, and timeliness. moreover, it does not have support for centralized filtering of events.The CORBA Notification Service was specified by the OMG to address the limitations with the CORBA Event Service.

2 Structure and Functionality in the CORBA Notification Service

This section first describes the structure of the core components[1] in the CORBA Notification Service [5].

2.1 Component Structure of the CORBA Notification Service

Figure 1 illustrates the components in the standard CORBA Notification Service. This architecture is similar to the architecture of the CORBA Event Service [2], though some components have a broader range of capabilities in the Notification Service. The enhancements in the Notification Service architecture are backwardly compatible to preserve interoperability with clients written for the CORBA Event Service. Each of these components is described below.

Structured Events: Structured events define a standard data structure into which a wide variety of event messages can be stored. The schema for structured events is known to the Notification Service and its clients. Consumers can install different filters that use the "filterable body" fields of the structured event definition to match with the filter constraint expressions efficiently. As shown in Figure 2, the header of a structured

*This work was funded in part by ATD, Cisco, Siemens MED, and DARPA ITO

[1] The term *component* used throughout this paper refers to a "component" in the general sense, *i.e.*, an identifiable entity in a program, rather than in the more specific sense of the CORBA Component Model [4].

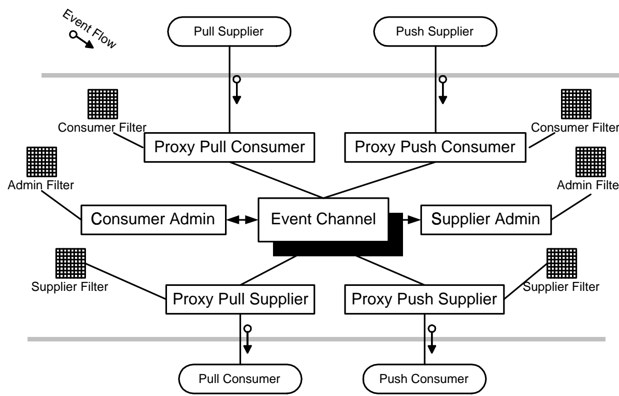

Figure 1: Components in the CORBA Notification Service

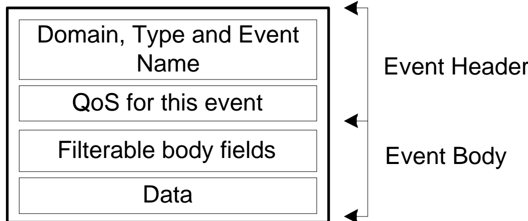

Figure 2: Structured Event

event consists of type information and a variable header, which can carry the QoS properties of an event. The event body consists of filterable body fields followed by the payload data.

Proxy objects: Proxy objects are delegates that provide complementary interfaces to clients, *i.e.*, a consumer obtains and connects to a proxy supplier and a supplier obtains and connects to a proxy consumer. Hence, a supplier sends events to its proxy consumer, whereas a consumer receives events from its proxy supplier. This abstraction enables anonymous connectivity between consumers and suppliers.

Admin objects: Each admin object is a factory [3] that creates the proxy interface to which each client will ultimately connect. Consumer admins create proxy suppliers to which consumers connect. Conversely, supplier admins create proxy consumers to which the suppliers connect. Supporting multiple admin objects in a given event channel enables the logical grouping of the proxy objects associated with the channel according to common subscription information. This feature is particularly useful with respect to consumer admin objects, since it enables the channel to optimize the servicing of a group of consumers that are interested in receiving the same set of events.

Filter and Mapping Filter Objects: Filter objects can be associated with all admin and proxy objects. Filter objects that affect the event forwarding decisions made by proxy objects encapsulate a set of constraints. Each constraint consists of (1) a sequence of event types and (2) a string containing a boolean expression whose syntax conforms to a constraint grammar. The default constraint language defined by the Notification Service is Extended TCL, which extends the TCL (Trader Constraint Language) specified by the Trading Service [6].

Event channels: An event channel is a factory that creates consumer admin and supplier admin objects. This differs slightly from the CORBA Event Service event channels, which only have one instance of admin objects. QoS and admin properties can be set on the event channel during its creation. These parameters are passed as default values to any admin object created by the channel. These parameters can be changed subsequently by consumers and suppliers.

QoS properties: The specification uses properties, *i.e.*, `<String, Any>` tuples, to define QoS properties. QoS properties can be associated with an event channel, admin objects, proxy suppliers and consumers, and individual event messages. The properties defined by the specification are:

- *Reliability* – The event reliability and connection reliability specify fault tolerance properties to the Notification Service. If these properties are supported then after a Notification Service is restarted after a crash, it must reconnect to all its clients and deliver all events that have not expired yet to its consumers.
- *Priority* – This property controls the order in which events are delivered to consumers. The event channel will attempt to deliver messages to consumers in priority order.
- *Expiration times* – This property indicates the time range in which an event is valid. If an event is not delivered within a specified time then an event channel should discard it.
- *Earliest delivery time* – This property specifies how long an event must be held in the channel before it is delivered.
- *MaximumEventsPerConsumer* – This property bounds the maximum number of events the channel will queue on behalf of a given consumer. This property helps avoid the case when the channel fills up its queues with events destined for a misbehaving consumer.
- *Order Policy* – This property specifies the order in which events are buffered for delivery.
- *Discard Policy* – This property specifies policies for discarding events when the queues are full.

Admin properties: The following administrative properties can be set on an event channel:

- *MaxQueueLength* – This property specifies the maximum number of events that will be queued by the channel before the channel begins discarding events or rejecting new events upon receipt of each new event.
- *MaxConsumers* – This property specifies the maximum number of consumers that can be connected to the channel at any given time.
- *MaxSuppliers* – This property specifies the maximum number of suppliers that can be connected to the channel at any given time.

3 Designing and Optimizing TAO's Notification Service

This section describes how we implemented the CORBA Notification Service in TAO [7], which is a CORBA-compliant ORB that supports applications with stringent QoS requirements. TAO's Notification Service makes it easier to develop distributed applications in heterogeneous environments by providing application transparency, high flexibility, scalability, interoperability, bounded resource consumption, filtering of events, and FIFO/deadline/priority-based event delivery.

The following design challenges were identified prior to and during the development of TAO's Notification Service:

1. Handling multiple event, supplier and consumer types uniformly.
2. Efficiently propagating different event types to different types of consumers.
3. Minimizing interference between the event channel participants.
4. Ensuring fairness in event processing.
5. Optimizing the performance of the CORBA `Any` type.
6. Optimizations for footprint reduction.
7. Customizing event channels for particular deployment environments.
8. Optimizing event filter evaluation.

These challenges and our solutions are discussed below. To enhance the generality of our solutions, we describe them in terms of the patterns [3, 8] we used to resolve the design challenges.

3.1 Challenge 1: Handling Multiple Event, Supplier and Consumer Types Uniformly

Context: Events transmitted between event channel participants can have different representation, such as `Any`s, structured events, and sequences of structured events.

Problem: An event channel must propagate events from suppliers that feed its events in any form to consumers that want the event in any other form. When events are of the same type, however, we do not want to perform needless conversions to a canonical format, nor do we want to copy the event multiple times in memory. Similarly, there are different types of consumers and suppliers that can connect to the event channel. Since IDL interfaces are similar we do not want to write a specific implementation for each type with redundant code in each one of them.

Solution → the Adapter pattern: This pattern converts the interface of a class into another interface that clients expect [3]. The Adapter pattern lets classes work together that could not otherwise due to incompatible interfaces. This pattern relies on object composition, *i.e.*, the adaptee object is contained by the adapter object. The adapter also implements a target interface, which is the interface expected by a client class. The client deals with and invokes methods, only on the target interface. The adapter implementation of the target interface delegates operations to the adaptee. Hence, various adaptees can conform to the target interface and thus maintain a single interface to a client.

Applying the Adapter pattern in TAO: Figure 3 shows how TAO uses a base class to represent an event. The `Any_Event` and `Structured_Event` classes (adapters) "adapt" the `CORBA::Any` and `CosNotification::StructuredEvent` (adaptees) to the `Notify_Event` interface (target). These classes know how to manage memory and convert between the various types, *i.e.*, `Any` and `Structured`. This design results in a uniform treatment of events throughout TAO's Notification Service event channel implementation, thereby minimizing

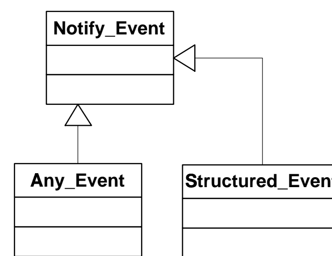

Figure 3: The `Notify_Event` Abstraction

code duplication and allowing the integration of different strategies to process all types of events.

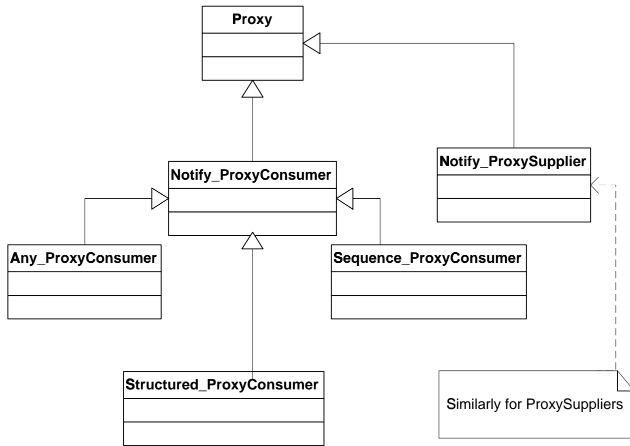

Figure 4: Abstraction for Proxy Objects

Figure 4 shows the classes TAO's Notification Service uses to represent proxy objects. The `Any_ProxyConsumer`, `Structured_ProxyConsumer` and `Sequence_ProxyConsumer` classes (Adapters) are used to adapt the `CosEvent::ProxyConsumer`, `CosNotifyAdmin::StructuredProxyConsumer`, and `CosNotifyAdmin::SequenceProxyConsumer` interfaces (Adaptees) to the `Notify_ProxyConsumer` (Target) interface. Similarly, the `ProxySupplier` classes are adapted to the `Notify_ProxySupplier` interface.

3.2 Challenge 2: Efficiently Propagating Different Event Types to Different Types of Consumers

Context: When an event is send from a supplier to the event channel, the receiving consumer(s) might not accept the same type of event, *e.g.*, an event send as an `Any` type could be received as an `Any`, a structured event, or a sequence of structured events.

Problem: We need to propagate different event types to different types of consumers efficiently.

Solution → the Visitor pattern: This pattern decouples operations that traverse elements in a complex object structure from the object structure itself [3]. The Visitor pattern lets us define a new operation without changing the classes of the elements on which it operates. When a visitor object calls the `accept` method of an element in an object structure, the implementation of the `accept` operation calls back the visitor object's `visit` method and passes information about its own concrete type.

Applying the Visitor pattern to TAO: Figure 5 illustrates how each consumer type implements the `dispatch_event` method, which in turn invokes the `push_event` method for consumer-specific event dispatching. When an event is selected to be dispatched to a consumer, the event processing engine invokes the `dispatch_event` method of the proxy consumer. In turn, the specific implementation of this virtual method invokes the correct `push_event` method of the event. This method then performs any necessary type conversion and initiates event dispatching to the remote consumer.

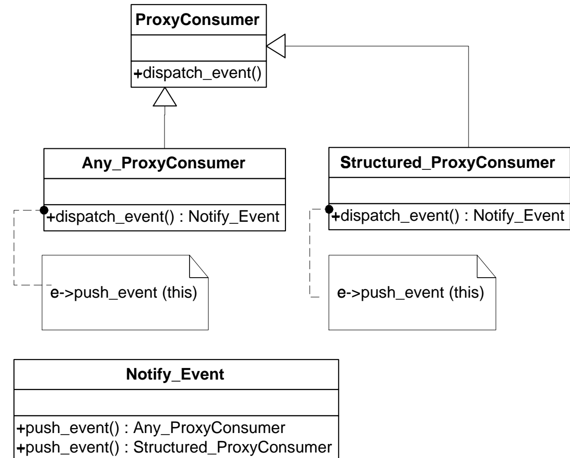

Figure 5: Double Dispatching for Event Propagation

3.3 Challenge 3: Minimizing Interference Between Event Channel Participants

Context: An Event channel should be able to handle event delivery from suppliers and should be able to perform event forwarding with the minimum possible latency, *i.e.*, suppliers delivering an event to the channel should not have to wait while the channel forwards events to recipient consumers. Similarly when the event channel forwards events to multiple consumers, each consumer might spend an unbounded amount of time in the implementation of its `push` method.

Since events are forwarded by the event channel via a CORBA two-way method, the channel has no choice but to have the dispatching thread wait while the consumer returns from the call. Another case of such a coupling occurs when a method on a filter object is invoked to check if an event's properties match the filter constraints. A constraint could be arbitrarily complex and the filter itself could be a remote object. These factors can affect the event channel's event processing time.

Problem: A reasonable implementation should strive to minimize the interference between different participants of the

event channel.

Solution → the Active Object pattern: This pattern decouples method execution from method invocation in order to simplify synchronized access to an object that resides in its own thread of control [8]. The Active Object pattern allows one or more independent threads of execution to interleave their access to data modeled as a single object.

Applying the Active Object pattern to TAO: Using the Active Object pattern at the various stages of event processing enables the minimization of the interference between the event channel participants. As shown in Figure 6, events and the operation performed on them are encapsulated as command

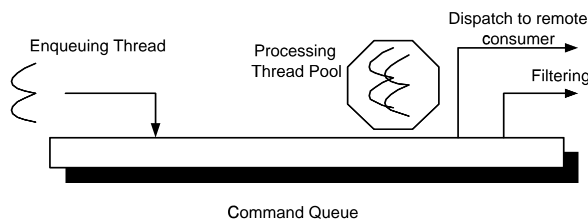

Figure 6: Asynchronous Event Processing Using the Active Object Pattern

objects. Enqueueing thread(s) place events into a command queue according to a buffering order policy. An Active object with worker threads dequeues and executes the command objects in the queue. Note that the ORB itself can be configured to increase concurrency by using the Leader/Followers pattern [8]. In this case, each ORB invocation is handled by a separate thread, allowing multiple events to be delivered concurrently to the event channel.

3.4 Challenge 4: Ensuring Fairness in Event Processing

Context: The priority QoS property specifies the relative importance of an event. The Notification Service ensures that priorities are respected by enqueueing events in an internal buffer according to priority.

Problem: A long-duration filter evaluation operation involving a maximum of four remote filters (1 each for the Proxy Objects and 1 each at the Admin Objects) can starve other events in the queue and prevent them from being processed promptly.

Solution → the Command Object pattern: This pattern encapsulates a request as an object, thereby allowing parameterization of different requests [3]. The actual nature of the request is hidden by the Command Object interface. Different concrete implementation of the Command interface implement the request and provide semantics to it. This pattern can be used to decompose the internal event processing within an event channel into stages to ensure fairness.

Applying the solution in TAO: The filter evaluation, subscription lookup and event dispatching operations are encapsulated as command objects. As shown in Figure 7 instead of performing these six operations synchronously, the evaluation

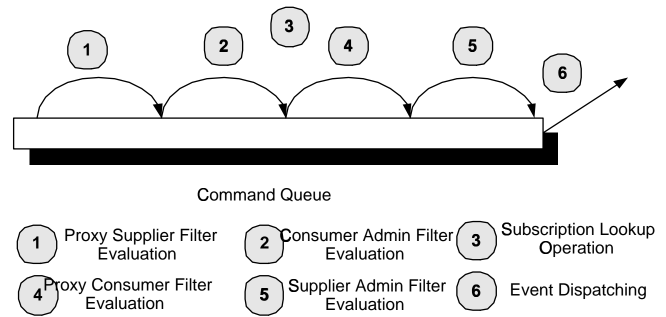

Figure 7: Processing Command Objects to Ensure Fairness

is broken up into discrete operations. If the event is still eligible for further processing after executing an command operation, it is enqueued back into the command queue as the following command object and subsequently dequeued for further processing. Finally, the dispatching command sends out the event to the remote consumer.

3.5 Challenge 5: Optimizing the Performance of `Any`s

Context: A CORBA-compliant Notification Service must be able to process events containing `Any`s.

Problem: In CORBA, `Any`s are expensive data types because they can have many levels of nesting. For example, an `Any` can be contained within a structure, which can itself be contained within an `Any` and so forth. When a demarshaling engine decodes this expensive type from a common data representation (CDR) stream, it makes a copy of the entire data buffer used to represent the `Any`. Likewise, copying an `Any` can require several memory allocations and buffer copys to obtain a new representation of the CDR stream. Moreover, the C++ mapping of CORBA `Any`s requires them to be responsible for any memory returned to the application. Optimized ORBs should share the `Any` contents even if there are multiple copies of the `Any` object.

Solution → Reference counting via the Handle/Body idiom: This presents multiple logical copies of the same data while sharing the same physical copy [9]. In C++, this idiom is often used to automate the memory management in conjunction with reference counting and smart pointers.

Applying the solution in TAO: In TAO the CDR marshaling engine does not copy the CDR stream into the `Any`, instead, all CDR streams are reference counted, and the `Any` only increments the reference count to maintain a logical copy of the buffer. Likewise, once the contents of the `Any` are extracted by the application the `Any` object becomes responsible for deallocating the extracted object. This extracted object can be shared by multiple instances of the `CORBA::Any` object, minimizing the cost of copying and extracting the contents repeatedly. The use of the Handle/Body idiom implements this optimization without changing the semantics required by the standard C++ mapping.

3.6 Challenge 6: Optimizations for Footprint Reduction

Context: The Notification Service specification has many features that might be required by all applications, *e.g.*, some deeply embedded systems may not want to incur the increase in memory footprint for certain unneeded features.

Problem: A required set of Notification features should be "composable" by users.

Solution → the Builder pattern: This pattern separates the construction of a complex object from its representation so that the same construction process can create different representations [3].

Applying the solution in TAO: The Builder pattern uses the appropriate sets of libraries to compose a configuration required in a particular use-case. For example, a configuration containing *no-filtering + AnyProxySupplier + AnyProxyConsumer + reactive dispatching strategy* would yield the semantics of the CORBA Event Service.

These features are separated into libraries as follows:

- The three different pairs of proxy supplier and proxy consumer types are separated into different libraries. In a specific configuration, only the required type (*e.g.*, the push proxy supplier) may be loaded by a builder at startup. Hence, the other types of proxy implementations are not loaded since they are not needed.
- An application may not require filtering, in which case the filtering engine library is not loaded by the builder.
- The Dispatching Strategies could be simple reactive dispatching, or asynchronous dispatching as described in Section 3.2.

3.7 Challenge 7: Customizing Event Channels for Particular Deployment Environments

Context: The standard CORBA Notification Service is configurable in the following manner:

1. A user can specify features required by configuration.
2. A specific class implementation can be modified by the user to enhance or customize behavior.
3. Users can vary default properties, such as thread pool size and locking strategy.

Problem: A mechanism is needed to allow the application developers to use the various configurable option in the Notification Service.

Solution → the Component Configurator pattern: This pattern decouples the behavior of component services from the point in time at which service implementation are configured into an application [8].

Applying the solution in TAO: All objects in TAO's Notification Service implementation are created via factory objects. These factories can be loaded statically or dynamically by using the ACE framework [10]. Figure 8 shows a schematic of the how the Component Configurator is used in the Notification Service.

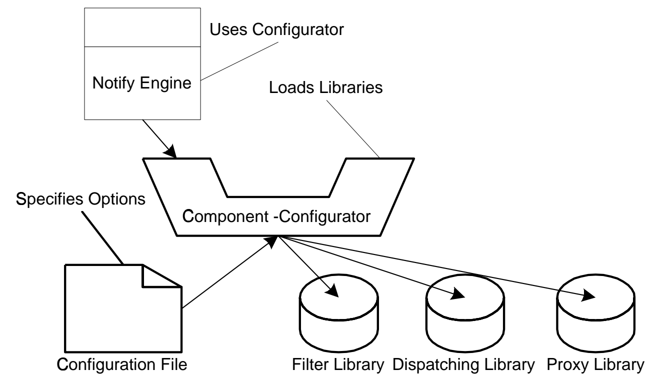

Figure 8: Apply the Component Configurator Pattern in the Notification Service

3.8 Challenge 8: Optimizing Filter Evaluation

Context: Admin objects in the Notification Service can be associated with filter objects. All proxy's connected to such an admin share the list of filter objects associated with the admin. Moreover, the proxys themselves are associated with filter objects. An event must satisfy the constraints specified by the filters at both the proxy and admin level.

Problem: After an admin level filter has been matched successfully against an event, we do not want to repeat this matching operation for each proxy connected to the same admin.

Solution → Optimization principle pattern of "passing hints": We use a variation of the optimization principle pattern "passing information between layers" [11]. This pattern is commonly used in protocol stack optimizations where each protocol layer passes certain information to the layer on top to help it avoid demultiplexing overhead.

Applying the solution in TAO: In TAO's Notification Service, we pass a hint to proxy objects to skip filter evaluation of their parent admin object if this has already been performed. We can optimize the filtering of a given event by a group of proxies since each member of the group logically applies the same filters to the same event. Thus, the results of the evaluation of a given event against a given filter can be shared by all proxy objects managed by a given admin object.

Figure 9 shows the configuration of a consumer admin (with

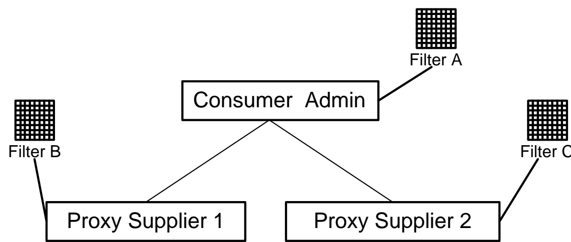

Figure 9: The Optimizing Filter Evaluation

filter `A`) of proxy supplier 1 (with filter `B`) and proxy supplier 2 (with filter `C`). An event must pass filters at both levels to be dispatched to consumers[2]. Filter `A` is evaluated first and the result of that evaluation is passed as a hint to both the proxy suppliers. This hint is used subsequently to determine if a proxy filter should be evaluated.

4 Empirical Benchmarking Results

This section presents the results of empirical benchmarks that quantify various aspects of TAO's Notification Service performance.

4.1 Overview of the Testbed Environment and Benchmarks

The experiments were conducted on a QuadCPU PC with 400MHz Pentium2 processors, 512KB on-chip cache, with 256MB of free RAM memory and running the Linux operating system version 2.2.18. The benchmarking programs were compiled using the GCC compiler version 2.95.4, with all optimizations enabled. All benchmarks were run in the POSIX real-time thread scheduling class [12].This scheduling class enhances the integrity of our results by ensuring the threads created during the experiment were not preempted arbitrarily during their execution. Using this testbed, we conducted the following performance experiments:

- *Event channel scalability* – In this experiment we show the behavior of the Event Channel with progressively increasing number of consumers with each consumer receiving the same event from one supplier transmitting a burst of events.
- *Measuring the effect of Thread configuration* – In this experiment, we measure the throughput of consumers, obtained in different thread configuration scenarios.
- *Event channel filtering overhead* – In this experiment, we measure the impact of filtering events on channel throughput by setting a filter at the consumer admin and measuring the throughput obtained before and after setting the filter.

4.2 Measuring Event Channel Scalability

Scalability measurement technique: In this experiment, we measure the scalability of TAO's Notification Service, where scalability is defined in terms of per-consumer throughput as the number of event consumers increase. In this test there is no event filtering *i.e.*, every consumer receives each event sent to an event channel.

Within an event channel we use a thread pool to process incoming event messages. The pool size is increased from 1 (reactive) to n to measure how the increase in concurrency affects performance.

Scalability measurement results: The per-consumer throughput is measured with 1, 5, 15 and 20 consumers, and 1 supplier.A burst of 1000 structured events is delivered with increasing number of dispatching threads in the Notification Service - none, 1, 2, 3, 4. All the consumers and suppliers are colocated in the same process which uses 1 worker thread to run the ORB event loop. The Notification Service is run in a seperate process on the same machine. Hence the two ORB's use the loopback interface to communicate, thereby eliminating any variance in the results due to network load.

figure 10 shows that the per-consumer throughput increases with increasing concurrency, however this speedup minimizes as the number of consumers increase, due to the overhead of synchronization and context-switching between the threads.

[2] An inter-filter group operator specifies if the results of evaluating the proxy and admin filters should be logically AND or OR

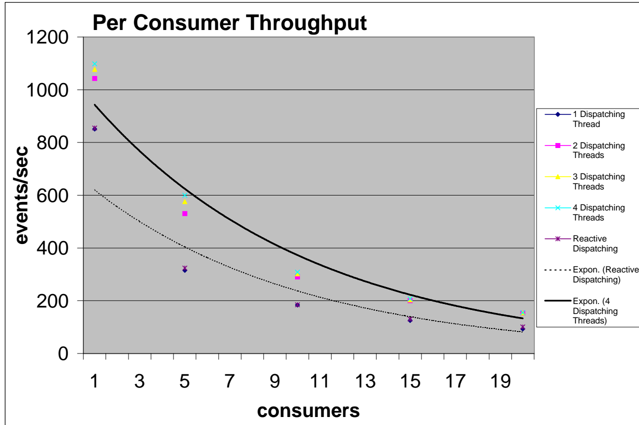

Figure 10: Consumer Throughput

4.3 Measuring the effect of Thread configuration

Thread configuration measurement technique: In this experiment, we connect 1 supplier and 2 consumers to the event channel. One of the consumers (slow) has a delay of 1 sec in it's *push* implementation. The other consumer (fast) has no delay.By configuring the Notification Service correctly, we will show how the throughput of the fast consumer can remain unaffected though it is connected to the same event channel as the slow consumer.

Case 1=Reactive, Case 2 = 2 Threads, 1 Shared Queue, Case 3 = Thread and Queue, per Proxy Supplier

	Consumer	Throughput (events/sec)
Case 1	Slow	0.99
	Fast	0.99
Case 2	Slow	0.99
	Fast	1.2
Case 3	Slow	0.99
	Fast	519

Table 1: Effect of Thread configuration on Throughput

Thread configuration measurement results: The per-consumer throughput is measured in 3 different configurations of the Notification Service. In the reactive case (*Case1*), there are no dispatching threads. In the second configuration (*Case 2*), events are queued in a single queue. This queue is accessed by 2 dispatching threads.In this case we notice a slight increase in the throughput of the fast consumer. In the third configuration (*Case 3*), we use a queue per proxy supplier with a dedicated thread that dequeues the events.Events destined for each consumer type are queued in their respective proxy suppliers. This configuration decouples the data path of each consumer, which greatly increases the throughput of the fast consumer.

4.4 Measuring Event Filtering Overhead

Filtering overhead measurement techniques: We send an event burst of 1000 structured events between 1 supplier and 1 consumer colocated in the same process. The Notification Service is run on the same host with 2 and 4 dispatching threads. The consumer throughput is measured without and then with a filtering constraint at the consumer admin.The measurements are obtained for 5, 15, 30 and 50 consumers.

NF=No Filter, F=Filter, 2T=2 Threads, 4T=4 Threads

		Throughput (events/sec)	
Consumers	**NF, 2T**	**F,2T**	**F,4T**
5	502	432	442
15	199	166	170
30	102	86	87
50	61.5	52	53

Table 2: Filtering Overhead

Filtering overhead results: By introducing a filter, the average throughput drops by 15% percent.By increasing the number of dispatching threads to 2, the throughput increases slightly by 2% for the 5 and 15 consumers case.

5 Related Work

A number of research projects have focused on distributed publish/subscribe mechanisms. For example, Rajkumar *et al.*, describe a real-time publisher/subscriber prototype developed at CMU SEI [13]. Their Publisher/Subscriber model is functionally similar to the CORBA Event and Notification Services, though it uses real-time threads to prevent priority inversion within the communication framework. One interesting aspect of the CMU model is the separation of priorities for subscription and event transfer so that these activities can be handled by different threads with different priorities. However, the model does not utilize any QoS specifications from publishers (suppliers) or subscribers (consumers). As a result, the message delivery mechanism does not assign thread priorities according to the priorities of publishers or subscribers. In contrast, TAO's Real-time Event Service [2] utilizes QoS parameters from suppliers and consumers to guarantee the event delivery semantics determined by a real-time scheduling service.

COBEA [14] is a CORBA-based event architecture service that generates parameterized events, which are published by a trading service. For scalability, clients must register their interest at the service, at which point an access control check is carried out. Subsequently, whenever a matching event occurs, the client is notified. This COBEA project is similar to the TAO Notification Service, the main difference being that TAO's Notification Service is based on the OMG standard.

There are several commercial CORBA-compliant Notification Service implementations available from vendors, such as Iona, Inprise, SunSoft [15], and DTSC. Iona also sells OrbixTalk, which is a messaging technology based on IP multicast. Unfortunately, since the CORBA Notification Service specification does not address issues critical for real-time applications, these implementations are not acceptable solutions for many domains.

6 Concluding Remarks

Many distributed applications, such as real-time avionics mission computing systems, distributed interactive simulation, and computer-assisted stock trading, require an event-based communication model. By using the Notification Service described in this paper, these applications can be built effectively by leveraging a middleware solution that is standards-based, flexible, and optimized for high-performance and scalability. The CORBA Notification Service builds upon the CORBA Event Service, which delivers events to all consumers connected to it on a best-effort basis. The Notification Service extends this service by providing the following two general mechanisms:

- *Event filtering*, which allow applications to control which supplier events are disseminated to which consumers. This selective control helps reduce redundant network traffic, *i.e.*, events that would be rejected by consumer applications after traversing the network are not sent in the first place.
- *QoS properties*, such as reliability, priority, ordering, and timeliness, which allow applications to bound resource consumption of the Notification Service. It also enables applications to specify the ordering of events in an event channel, thereby allowing events to be propagated with priorities and deadline criteria, rather than the strict FIFO ordering of the standard CORBA Event Service.

TAO's implementation of the Notification Service addresses the issue of scalability and high-performance by applying suitable patterns and reusable framework components, optimizing the critical path of event propagation at the service- and ORB-levels, and providing configurability to reduce memory footprint and customize deployment.

All the source code, documentation, examples, and tests for TAO and its Notification Service and Real-time Event Service mechanisms are open-source and can be downloaded from `www.cs.wustl.edu/~schmidt/TAO.html`.

References

[1] Object Management Group, *The Common Object Request Broker: Architecture and Specification*, 2.4 ed., Oct. 2000.

[2] 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), ACM, October 1997.

[3] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*. Reading, MA: Addison-Wesley, 1995.

[4] BEA Systems, *et al.*, *CORBA Component Model Joint Revised Submission*. Object Management Group, OMG Document orbos/99-07-01 ed., July 1999.

[5] Object Management Group, *Notification Service Specification*, OMG Document telecom/99-07-01 ed., July 1999.

[6] Object Management Group, *Trading ObjectService Specification*, 1.0 ed., Mar. 1997.

[7] 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.

[8] D. C. Schmidt, M. Stal, H. Rohnert, and F. Buschmann, *Pattern-Oriented Software Architecture: Patterns for Concurrent and Networked Objects, Volume 2*. New York, NY: Wiley & Sons, 2000.

[9] Jim Coplien, *Advanced C++ – Programming Styles and Idioms*. Addison-Wesley, 1992.

[10] D. C. Schmidt and T. Suda, "An Object-Oriented Framework for Dynamically Configuring Extensible Distributed Communication Systems," *IEE/BCS Distributed Systems Engineering Journal (Special Issue on Configurable Distributed Systems)*, vol. 2, pp. 280–293, December 1994.

[11] I. Pyarali, C. O'Ryan, D. C. Schmidt, N. Wang, V. Kachroo, and A. Gokhale, "Applying Optimization Patterns to the Design of Real-time ORBs," in *Proceedings of the* 5^{th} *Conference on Object-Oriented Technologies and Systems*, (San Diego, CA), USENIX, May 1999.

[12] Khanna, S., *et al.*, "Realtime Scheduling in SunOS 5.0," in *Proceedings of the USENIX Winter Conference*, pp. 375–390, USENIX Association, 1992.

[13] 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.

[14] C. Ma and J. Bacon, "COBEA: A CORBA-Based Event Architecture," in *Proceedings of the* 4^{rd} *Conference on Object-Oriented Technologies and Systems*, USENIX, Apr. 1998.

[15] Y. Aahlad, B. Martin, M. Marathe, and C. Lee, "Asynchronous Notification Among Distributed Objects," in *Proceedings of the* 2^{nd} *Conference on Object-Oriented Technologies and Systems*, (Toronto, Canada), USENIX, June 1996.