



Calhoun: The NPS Institutional Archive
DSpace Repository

NPS Scholarship

Theses

2002-03

A highly adaptable generic event-based message channel design for loosely coupling software modules

Eryigit, Cihat

Monterey, California. Naval Postgraduate School

<https://hdl.handle.net/10945/6064>

Copyright is reserved by the copyright owner

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

NAVAL POSTGRADUATE SCHOOL

Monterey, California



THESIS

**A HIGHLY ADAPTABLE GENERIC EVENT-BASED
MESSAGE CHANNEL DESIGN FOR LOOSELY
COUPLING SOFTWARE MODULES**

by

Cihat Eryigit

March 2002

Thesis Advisor:
Second Reader:

Geoffrey Xie
Chris Eagle

Approved for public release; distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE March 2002	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE: A HIGHLY ADAPTABLE GENERIC EVENT-BASED MESSAGE CHANNEL DESIGN FOR LOOSELY COUPLING SOFTWARE MODULES			5. FUNDING NUMBERS
6. AUTHOR(S) Eryigit, Cihat			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) DARPA and NASA			10. SPONSORING/MONITORING AGENCY REPORT NUMBER G417
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE Statement A
13. ABSTRACT (maximum 200 words) Component-based software engineering is an emerging software development approach based on the fundamentals of object-oriented technology. This approach moves programmers' focus from component development to component assembly. Event-based programming is one of the techniques that can be used to assemble software components into applications. In this thesis, a new, generic, highly adaptable and flexible event channel has been designed and implemented. The main product is a Java utility package, called "channel package", which should help Java application developers create or enhance large systems using an event-based programming approach. The new channel design has several demonstrated performance advantages over existing event channel implementations. The flexibility and adaptability of the channel package has also been validated by a successful upgrade of the channel mechanism of the SAAM prototype system.			
14. SUBJECT TERMS Event Programming, Event Channel, The Server and Agent based Active network Management (SAAM) Architecture, Inter-object communication.			15. NUMBER OF PAGES 176
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited.

**A HIGHLY ADAPTABLE GENERIC EVENT-BASED MESSAGE CHANNEL
DESIGN FOR LOOSELY COUPLING SOFTWARE MODULES**

Cihat Eryigit
Lieutenant Junior Grade, Turkish Navy
B.S., Turkish Naval Academy, 1995

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

**NAVAL POSTGRADUATE SCHOOL
March 2002**

Author: Cihat Eryigit

Approved by: Geoffrey Xie, Advisor

C. Eagle, Second Reader

C. Eagle, Chairman
Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Component-based software engineering is an emerging software development approach based on the fundamentals of object-oriented technology. This approach moves programmers' focus from component development to component assembly. Event-based programming is one of the techniques that can be used to assemble software components into applications.

In this thesis, a new, generic, highly adaptable and flexible event channel has been designed and implemented. The main product is a Java utility package, called "channel package", which should help Java application developers create or enhance large systems using an event-based programming approach. The new channel design has several demonstrated performance advantages over existing event channel implementations. The flexibility and adaptability of the channel package has also been validated by a successful upgrade of the channel mechanism of the SAAM prototype system.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
	A. SAAM PROJECT AND SAAM CHANNEL MODEL.....	2
	B. SCOPE OF THIS THESIS.....	3
	C. ORGANIZATION OF THIS THESIS.....	4
II.	BACKGROUND.....	5
	A. JAVA EVENT MODEL.....	5
	1. Java 1.0 Event Model.....	5
	2. Java 1.1 Event Model.....	7
	B. CORBA EVENT MODEL.....	11
	1. Object Management Group (OMG).....	11
	2. Corba.....	12
	3. CORBA Event Service.....	13
	a. <i>The Push Model</i>	14
	b. <i>The Pull Model</i>	15
	c. <i>Mixing the Push and Pull Model in a Single System</i>	16
	d. <i>Types of Event Communication</i>	16
	4. CORBA Notification Service.....	17
	a. <i>Filter</i>	18
	b. <i>Quality of Service (QoS)</i>	18
	c. <i>Structured Events</i>	19
	5. Conclusion.....	20
	C. SAAM EVENT CHANNEL MODEL.....	21
	1. SAAM Channel.....	21
III.	CHANNEL PERFORMANCE METRICS AND EVALUATION OF EXISTING SAAM CHANNEL MODEL.....	25
	A. CHANNEL PERFORMANCE METRICS AND QUALITATIVE OBJECTITIVES.....	25
	1. Channel Throughput (R).....	25
	2. Channel Work Rate (W).....	26
	3. Channel Access Delay (D).....	26
	4. Event Talk Time (T_t).....	26
	5. Thread Count (C).....	27
	6. Adaptability of Channel.....	28
	7. Scalability of Channel.....	28
	8. Manageability of Channel.....	29
	9. Functional Flexibility.....	29
	10. Ease of Use.....	29
	a. <i>Easy Creation of Channel</i>	30
	b. <i>Meaningful Method Names</i>	30
	c. <i>Easy Configuration of Channel</i>	30

	<i>d.</i>	<i>Sufficient and Understandable Method Overriding</i>	30
	<i>e.</i>	<i>Conflicts Between Methods</i>	30
	<i>f.</i>	<i>Conflicts Between Parameters in a Method</i>	30
	<i>g.</i>	<i>Documentation</i>	31
B.		TEST OF SAAM CHANNEL	31
	1.	Channel Throughput and Work Rate	31
	2.	Channel Access Delay	32
	3.	Event Talk Time	34
	4.	Thread Count	36
	5.	Manageability and Scalability of SAAM Channel	36
	6.	Adaptability and Functional Flexibility of SAAM Channel	38
	7.	Ease of Use of SAAM Channel	38
IV.		NEW CHANNEL MODEL AND CHANNEL PACKAGE	39
A.		INTRODUCTION	39
B.		FEATURES OF NEW CHANNEL DESIGN	39
	1.	Generic Event Structure	39
	2.	Event Buffering	40
	3.	Event and Channel Participant Priority	40
	4.	Event Delivery Order	40
	5.	Event Filtering	40
	6.	Self-Dispatching	40
	7.	Duplex Communication	40
	8.	Concatenating Channels	41
C.		CHANNEL PACKAGE	41
	1.	ChannelEvent Class	41
	2.	ChannelListener Interface	42
	3.	ChannelFilter Interface	43
	4.	ChannelScheduler Interface	44
	5.	ChannelListenerItem Class	45
	6.	ChannelEventFilter Class	46
	7.	Schedulers	49
	<i>a.</i>	<i>FIFOScheduler Class</i>	49
	<i>b.</i>	<i>PerTalker_RR_Scheduler Class</i>	49
	<i>c.</i>	<i>PriorityScheduler Class</i>	49
	8.	Channel Class	50
	<i>a.</i>	<i>Creating a Channel Object</i>	50
	<i>b.</i>	<i>Adding and Removing Listeners to Channel</i>	50
	<i>c.</i>	<i>Adding and Removing Talkers to Channel</i>	51
	<i>d.</i>	<i>Talking on Channel</i>	52
	<i>e.</i>	<i>Event Dispatching</i>	52
	<i>f.</i>	<i>Event Filtering</i>	53
	<i>g.</i>	<i>Listener Self-Dispatching</i>	54
	<i>h.</i>	<i>Duplex Communication</i>	55
	<i>i.</i>	<i>Concatenating Channels</i>	55
	9.	ChannelAccessAuthority Interface	56

10.	ChannelManager Class	56
V.	TEST AND RESULTS	59
A.	EVALUATION OF NEW CHANNEL DESIGN	59
1.	Channel Throughput and Work Rate	59
2.	Channel Access Delay and Event Talk Time	61
3.	Scalability and Thread Count	62
4.	Adaptability and Functional Flexibility of New Channel	63
5.	Manageability	63
6.	Ease of Use	64
a.	<i>Easy Channel Management</i>	64
b.	<i>Easy Self-Dispatching</i>	64
B.	INTEGRATION OF CHANNEL PACKAGE AND SAAM PROTOTYPE.....	65
1.	Removing Obsolete Classes and Interfaces	65
2.	PermissionTableEntry Class.....	66
3.	Changes to ControlExecutive Class	67
4.	Reducing the Number of Channels	67
5.	Event Priorities.....	69
6.	Additions to SAAM GUI	70
VI.	CONCLUSION	71
A.	LESSONS LEARNED	71
1.	Programming with Threads	71
2.	Integration with SAAM Prototype	71
B.	FUTURE WORK	72
1.	Communication Between Distributed Applications	72
2.	Automatic Sense for Self-Dispatching	72
APPENDIX A.	CHANNEL PACKAGE.....	73
APPENDIX B.	THE ADDITIONAL CLASS FOR SAAM PROTOTYPE	151
LIST OF REFERENCES	155
INITIAL DISTRIBUTION LIST	157

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 2.1	Button Event Process in Java 1.0 Event Model	6
Figure 2.2	Event Life Cycles in Java 1.1	9
Figure 2.3	Java Event Class Hierarchy.....	11
Figure 2.4	ORB-to-ORB Communication.....	12
Figure 2.5	Push Model	15
Figure 2.6	Pull Model.....	15
Figure 2.7	Mixed Model.....	16
Figure 2.8	SAAM Channel Concept	22
Figure 3.1	Test Bed For Studying Impact of Listener Event Handling Time on Channel Throughput	31
Figure 3.2	SAAM Channel Throughput and Work Rate versus Listener Event Handling Time	32
Figure 3.3	Test Bed for Measuring SAAM Channel Access Delay.....	33
Figure 3.4	Average Access Delay versus Number of Channel Talkers	33
Figure 3.5	SAAM Channel Access Delay versus Event Handling Time	34
Figure 3.6	Event Talk Time Versus The Number of Talkers.....	35
Figure 3.7	Test Bed for Measuring SAAM Channel Event Talk Time	35
Figure 3.8	Event Talk Time Versus Event Handling Time.....	36
Figure 3.9	Event Concatenation in SAAM Channels.....	37
Figure 3.10	Two-way communication with SAAM channel	38
Figure 4.1	Channel Event Structure	42
Figure 4.2	Channel Scheduler	45
Figure 4.3	Channel Listener Item.....	45
Figure 4.4	Self-dispatching	46
Figure 4.5	ChannelEventFilter Filtering Process.....	48
Figure 4.6	Per Talker Round-Robin Scheduler	49
Figure 4.7	Channel Event Dispatching.....	53
Figure 4.8	Channel Event Filtering Process	54
Figure 4.9	Two-way Event Communication	55
Figure 4.10	Concatenating Channels.....	55
Figure 5.1	Test Bed For Studying Impact of Self-Dispatching on Channel Throughput and Work Rate.....	59
Figure 5.2a	Effect of Self-Dispatching on Channel Throughput	60
Figure 5.2b	Effect of Self-Dispatching on Channel Work Rate.....	60
Figure 5.3	Event Talk Time versus Number of Talkers.....	62
Figure 5.4	Channel Access Delay versus Number of Talkers.....	62
Figure 5.5	Encapsulating SAAM Events in ChannelEvent Structure	66
Figure 5.6	Necessary Channels for A New Interface on Existing SAAM Prototype	68
Figure 5.7	Snapshot of New Channel Debug Window	70

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 5.1	New Channel Structure in SAAM Prototype	69
-----------	---	----

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to Professor Geoffrey Xie and LCDR Chris Eagle for their guidance and help throughout this thesis research.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

Component oriented programming is an emerging technology based on the fundamentals of object-oriented technology. Component-oriented programming and design patterns have revolutionized software development. With the decomposition of monolithically built systems into independently alterable and extensible components, methodologies from other engineering disciplines are being successfully applied to software development. [Ref 1]

Assembly of components is becoming an increasingly common approach to development in the software industry. This is due in large part to its promise of considerably more productive programming. This increase in productivity results from a greater degree of independence between the development of one component and that of the other components, as well as individual component independence from the development of the larger overall assembly itself. Users of a component need only be concerned with its interface — the abstraction, or functionality, it provides — and need not be aware of the implementation underlying that function. Developers of a component need not be aware of the bigger picture of its use — their implementation simply needs to fulfill the functional obligations of the contract constituted by its interface. Different components can be developed by different teams, at different times, at different vendors or enterprises. [Ref 2]

Event-based programming is one of the techniques that can be used to assemble software components into applications. In this approach, components exchange information in an asynchronous manner. A component that produces output data does not communicate directly with the data consumer(s) to move the data. Instead, it calls up an independent delivery service (typically a system mechanism) for the data transport. The delivery service requires the data, along with pertinent information about its producer and consumer(s), to be encapsulated in generic container objects called events. A data-producing component is also referred to as an event source, a consumer is referred to as an event listener, and the delivery service as an event dispatcher. For ease of identification and handling, events are often classified based on functionality into

different types, with events of the same type tagged by the same unique identifier. Thus, a component can be an event source, a listener, or both, albeit for different event types.

Software components should be coupled as loosely as possible to maximize system modularity and to support dynamic update and replacement of components. However, in existing event handling models, event sources and event listeners are often tightly coupled. For example, the Java Event Delegation model requires that an event listener must be registered with a specific event source to be notified of events generated by that event source. This methodology results in a direct linkage between an event source and its listeners. As such, it is cumbersome to replace the event source, having to reestablish the linkage.

Another drawback of existing event models is that they are generally designed for communication between GUI components and event handlers. This causes limitations and problems when we try to adapt those event models to communication between software components other than GUI components.

A. SAAM PROJECT AND SAAM CHANNEL MODEL

The goal of the Server and Agent based Active network Management (SAAM) project is to generate a solution that will provide a guaranteed quality of service (QoS) while maintaining the relative simplicity and robustness of the underlying TCP/IP architecture. SAAM seeks to provide this QoS by introducing an additional network service that, when requested, can assist applications by reserving the necessary network resources. The SAAM architecture is designed to allow network engineers to incrementally replace existing internal infrastructure. [Ref 3]

A Java-based SAAM prototype has been built incrementally. The prototype consists of modular, self-aware, agent-based components. These components are loosely coupled, based on a novel event model called the SAAM Channel Model. In this model, the event delivery mechanism is implemented in the user domain and in the form of application-specific event channels. Each channel serves as a public communication medium for a particular set of components that perform a particular set of functions. This

association of channel with functionality eliminates the need for any direct linkage between event source and listener.

The current SAAM channel model was implemented by NPS graduates Dean Vrable and John Yarger in their thesis “The Server and Agent based Active network Management (SAAM) Architecture: Enabling Integrated Services”. The implementation provides event dispatching between SAAM components. The SAAM channel model offers several benefits over the standard Java Delegation Event Model by allowing delivery of events from one or more event sources (called talkers in the SAAM channel model) to one or more event listeners and by allowing event listeners to register with a channel that has no talkers. The details of SAAM Channel model will be described in Chapter 2.

However, the current SAAM Channel implementation has several major shortcomings. It is not a general-purpose utility as it supports only an application-specific event object. It has no event buffering capability and the performance is not optimized. For example, when an event is generated the thread of the event source is blocked unnecessarily until that event is delivered to all listeners. Further, the model does not support quality of service nor give a programmer a choice of dispatching events to multiple listeners in a specific order.

B. SCOPE OF THIS THESIS

The primary goal of this thesis is to identify and address the shortcomings of the current SAAM channel implementation. The thesis develops a highly configurable channel to loosely couple software components, following the event-based programming paradigm. This channel implementation supports multiple types of event objects and allows a programmer to choose different event dispatching schedules based on event priorities and the properties of the channel participants. Several performance metrics are defined for evaluating different event channel implementations, and the performance of the new channel implementation is optimized with respect to those metrics.

C. ORGANIZATION OF THIS THESIS

The remaining part of this thesis is organized into the following chapters.

- Chapter II: Background. Provides information about the Java Event Model, the CORBA Event and Notification Service, and the SAAM Channel Model.
- Chapter III: Channel Performance Metrics and Evaluation Of Existing SAAM Channel Implementation. Defines and explains the channel performance metrics used in this thesis and shows some experimental results on the performance of the existing SAAM channel implementation measured by these metrics.
- Chapter IV: New Channel Design and Java Channel Package. Describes the new channel design and implementation, and explains the classes and interfaces in the resulting Java channel package.
- Chapter V: Tests and Results. Presents results of performance tests of the new channel implementation and describes the integration of the new Java channel package into the existing SAAM prototype.
- Chapter VI: Conclusion. Summarizes the results from this thesis and outlines possible future work.

II. BACKGROUND

It is useful to consider some related event models before describing the new channel design. The Java event model, the CORBA event service, and the current SAAM channel model are described here.

A. JAVA EVENT MODEL

1. Java 1.0 Event Model

The Java 1.0 event model provided a methodology for processing events implemented through Java 1.0.2. It was superseded by the Java 1.1 event model with the release of JDK 1.1 in 1997.

Under the Java 1.0 event model, every type of event is encapsulated in a single class, the `Event` class, which is contained in the `java.awt` package. An event object contains information about the type of event it represents, when the event was generated, where the event occurred, and what keys were pressed when the event was generated.

In Java 1.0 event model, event types are limited and an event was delivered only to the component that generated it or one of its parent containers. The event is handed to the component's `handleEvent` method. The `handleEvent` method calls an event handler method based on the type of event it is processing. For example, if an action event is sent to the `handleEvent` method, the method will in turn call the `action` method. The event handler method will return a value of `true` if the event was completely processed or `false` if it was not. If the `handleEvent` cannot dispatch the event, the event is automatically forwarded to the component's container. If the container does not handle the event, it passes the event on to its parent container, and so on. In this way, events are propagated up the containment hierarchy until they are either consumed or reach the top level. If an event reaches the top level and a proper handler is still not found, the event will be discarded. Figure 2.1 illustrates how a button event would be processed with the Java 1.0 event model.

The system provides a default implementation for each of the event handling methods. The default implementations do nothing and return the value `false`. It is straightforward to override one or more of these methods to provide the desired event-processing methodology.

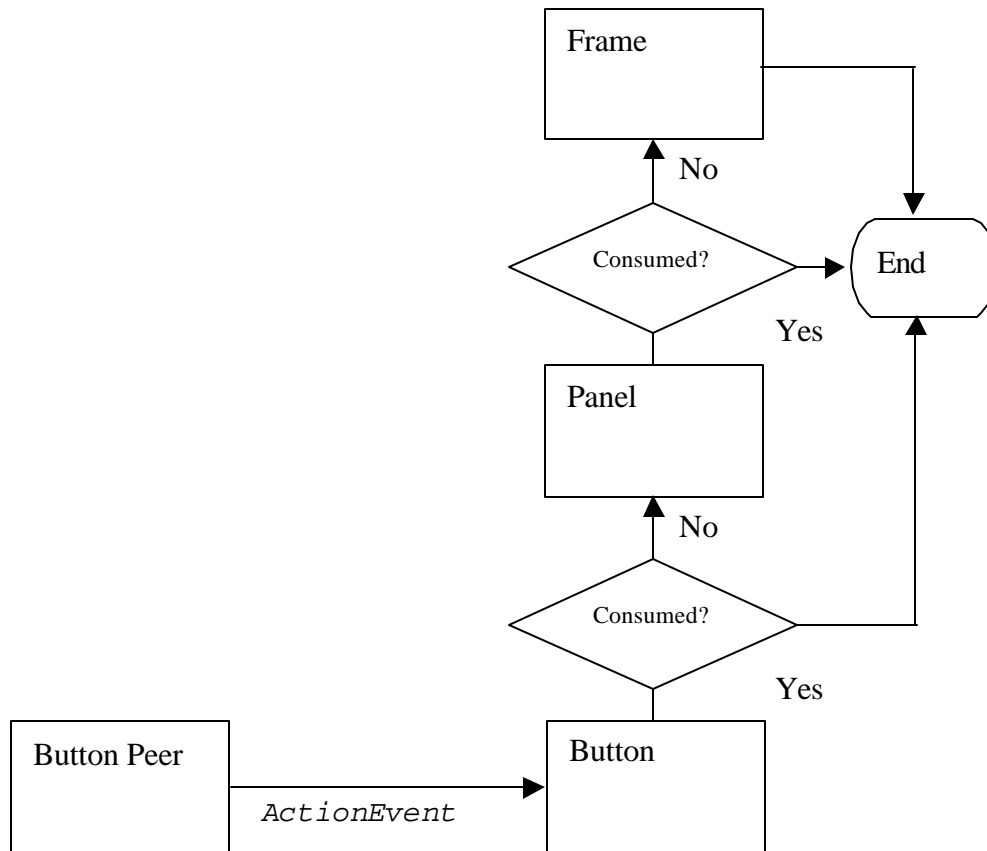


Figure 2.1 Button Event Process in Java 1.0 Event Model

One of the drawbacks of the Java 1.0 event model is that the event handler methods are defined in the `Component` class, and as such any `Component` subclasses can override these methods. This means the event handling tasks must be performed by the GUI components themselves. It is very difficult to decouple the event handling code from the GUI code.

Another drawback of this model is there is no filtering of events. Events are always delivered to the components regardless of whether the components actually

handle them or not. This can be a serious performance problem, particularly with high-frequency event types. [Ref 4]

There is also no way to use this event model for inter-object communication because the model restricts an event within the scope of one `Component` class object.

2. Java 1.1 Event Model

Deficiencies of the Java 1.0 event model became readily apparent. It wasn't very efficient to have the system search for the target component every time an event was generated. Another problem was the event handling code was explicitly tied to the GUI code creating inheritance problems. For example, instead of using one `Button` class for all button objects, an application might have to create a `Button` subclass that would quit when pressed, another `Button` subclass that would open a file when pressed, and so on.

One of the significant changes from Java 1.0 to 1.1 was the way events were handled. The Java 1.1 event model employs a concept called *delegation* and cleans up many deficiencies of the Java 1.0 event model. An event source generates the event and then *delegates* the event handling process to another piece of code. The event-handling object can be completely separate from the event source. Under this model any class can serve as the event handler. [Ref 4]

An event object is an instance of a subclass of `java.util.EventObject`; it holds information about the event. The `EventObject` class serves mainly to identify event objects; the only information it contains is a reference to the event source (the object that generated the event). [Ref 5]

In the Java 1.1 event model, an event source delivers an event only to registered listener objects. Listeners that are not registered with the event source do not receive that event. For each of the events an object can generate it maintains a list of listeners that are registered to receive the event. The object registers a listener by adding a reference to the listener to the proper listener list(s). This is done using one of the object's *add-listener* methods, passing the method a reference to the listener object as an argument. An object deregisters or disconnects from an event listener by removing the listener from its listener

list(s). This is accomplished by calling one of the object's *remove-listener* methods. [Ref 4]

An event is delivered by passing it as an argument to the receiving object's event handler method. `ActionEvents`, for example, are always delivered to a method called *actionPerformed* in the receiver. For each type of event, there is a corresponding listener interface that prescribes the method(s) an event handler must provide to receive the event. In this case, any object that receives `ActionEvents` must implement the `ActionListener` interface.

All listener interfaces are subinterfaces of the `java.util.EventListener`, which is an empty interface. It exists only to help the compiler identify listener interfaces. Listener interfaces are required for a number of reasons. First, they help to identify objects that are capable of receiving a given type of event. This way we can give the event handler methods friendly, descriptive names and still make it easy for documentation, tools, and humans to recognize them in a class. Second, listener interfaces are useful because several methods can be specified for an event listener. For example, the `FocusListener` interface contains two methods:

- `abstract void focusGained (FocusEvent event)`
- `abstract void focusLost (FocusEvent event)`

Although these methods both take a `FocusEvent` as an argument, they correspond to different reasons for firing the event. In the example, the reason is whether the `FocusEvent` received or lost the focus. Even though the event parameter passed to the methods contains enough information to determine whether the focus was gained or lost, requiring two methods, the `FocusListener` interface minimizes the effort. Specifically, if the *focusGained* method is called, it indicates that the event type was `FOCUS_GAINED`. [Ref 5]

Under the Java 1.0 event model, the dispatching and processing of events was a linear process. An event was sent to a single target component. If the event was not

completely processed by that component, it could be sent to another component, but the system was not able to broadcast an event to multiple event listeners simultaneously. Under the Java 1.1 event model, events can be sent to any number of event handler objects. Any listener class that is registered with the source component would receive the event.

Figure 2.2 shows the life cycle for events that are subclasses of the `AWTEvent` class. The `dispatchEvent` and `processEvent` methods take an `AWTEvent` object as an argument. However, many of the event classes contained in the `java.swing.event` package are not subclasses of `AWTEvent`, but instead they inherit directly from `EventObject`. The objects that generate these events will also define a `fireEvent` method that causes a given event to be delivered to all listeners in the event's listener list.

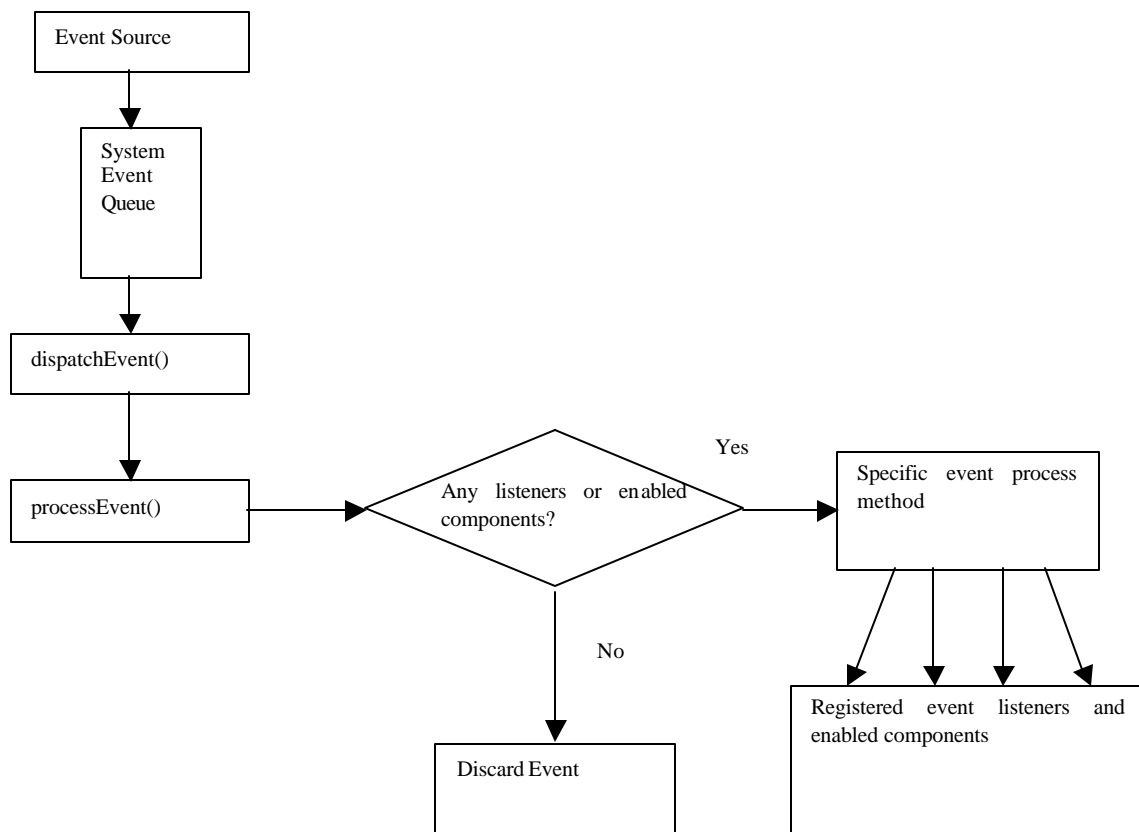


Figure 2.2 Event Life Cycles in Java 1.1

A system event queue is used to store newly generated events until they are dispatched by the *dispatchEvent* method. Normally, the operation of the system event queue is transparent to the user. It does what it does in the background, automatically. The `EventQueue` class encapsulates the Java event queue. It is possible to look into, and even manipulate, the system event queue via the `EventQueue` class.

All event-handling code executes in a unique thread, called the *event-dispatching thread*. It is the event-dispatching thread that calls any event listener methods. The event-dispatching thread retrieves and processes events from the system event queue in a First-In-First-Out (FIFO) fashion. It finishes executing an event handler's code before invoking the next event handler. One reason this is done is to keep the component in sync with the component display and to block other activities on a component while an event is being processed. For instance, when a button is pressed it will appear to sink into its container display and its color will change. While the `ActionEvent` that is generated is being processed, the event-dispatching thread will block any other user interactions with the button.

The Java API provides a large selection of event and listener classes. It also provides the basic building blocks for creating application-specific event classes and event listeners. One can define a new event class either from scratch, using the high-level event superclasses, `EventObject` and `AWTEvent`, or write a subclass of an existing `EventObject` or `AWTEvent` subclass. The event and support class hierarchy is shown in Figure 2.3 [Ref 4].

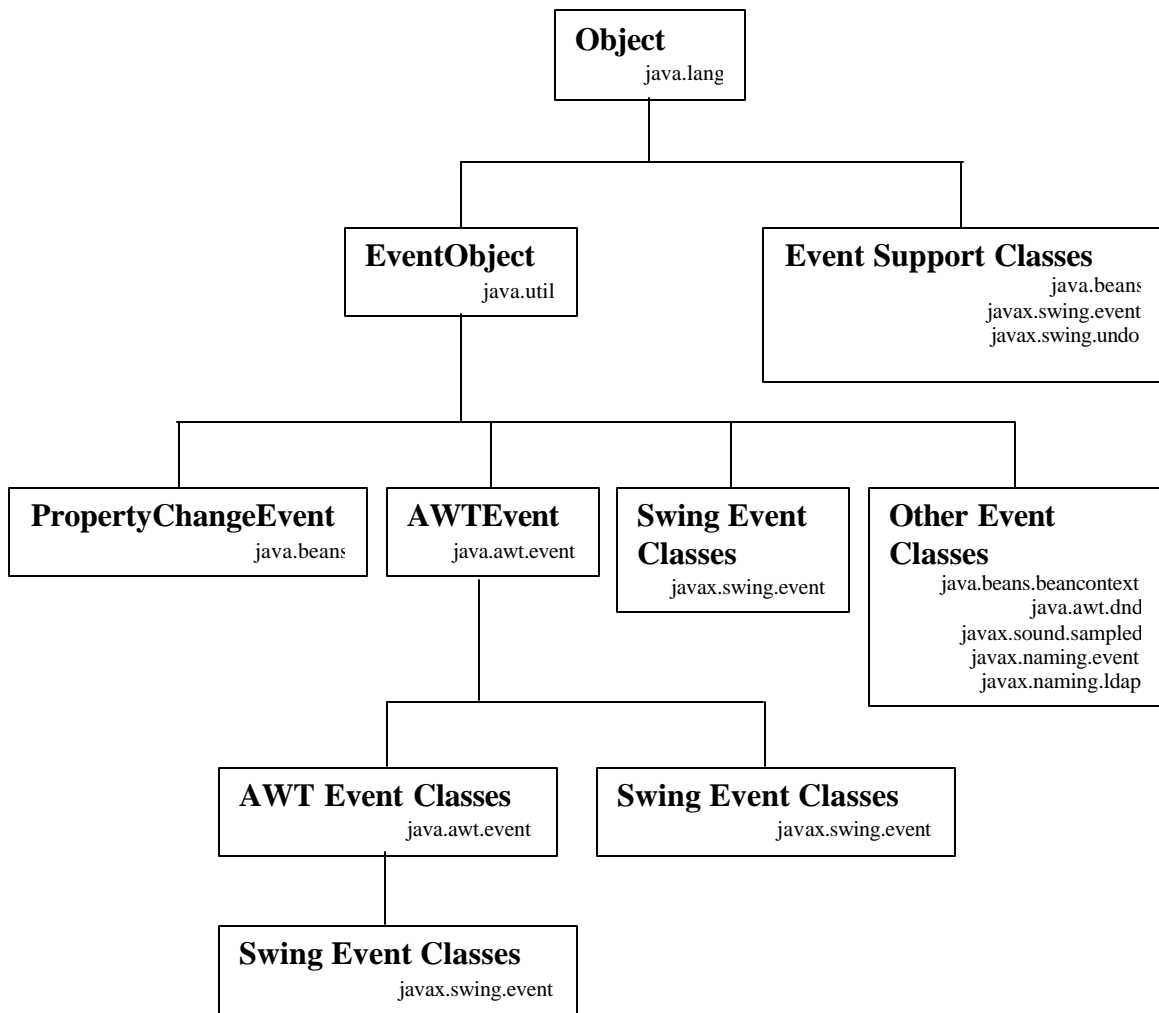


Figure 2.3 Java Event Class Hierarchy

B. CORBA EVENT MODEL

1. Object Management Group (OMG)

The Object Management Group is an international organization supported by over 800 members including information system vendors, software developers, and users. It was founded in 1989. OMG promotes the theory and practice of object-oriented technology in software development. The organization's charter includes the establishment of industry guidelines and object management specifications to provide a common framework for application development. The primary goal is to achieve high reusability, portability, and interoperability of object-based software in distributed,

heterogeneous environments. Conformance to these specifications makes it possible to develop a heterogeneous applications environment across all major hardware platforms and operating systems. [Ref 6]

2. Corba

CORBA is the acronym for Common Object Request Broker Architecture, OMG's open, vendor-independent architecture and infrastructure established to enable object-oriented computer applications to work together over networks. CORBA is the Object Management Group's answer to the need for interoperability among rapidly proliferating hardware and software products. It allows applications to communicate with one another no matter where they are located or who has designed them.

The CORBA specification only defines a set of conventions and protocols that must be followed by CORBA implementations. It is left to vendors and developers to translate this specification into a working implementation. CORBA does not make any restriction on language usage or underlying operating systems.

Because CORBA is language independent, it relies on an *Interface Definition Language (IDL)* to express how clients will make a request for a service.

A client communicates to a server object through an object reference. This is a pointer to the object that allows for operations and data access requests to be sent from the client to the server via an *Object Request Broker (ORB)*.

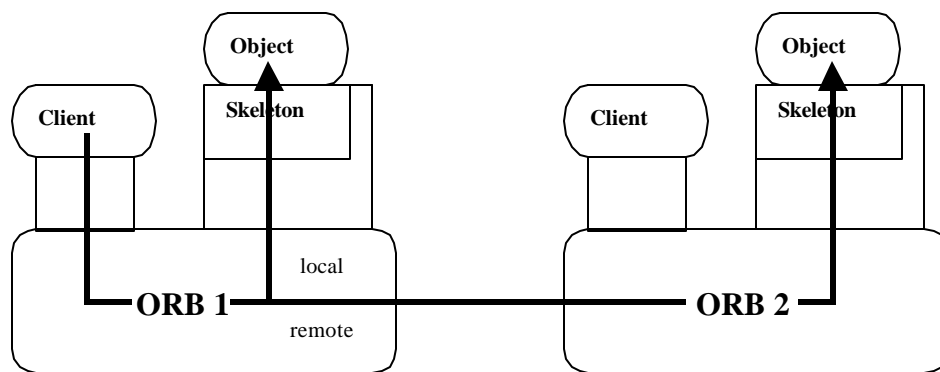


Figure 2.4 ORB-to-ORB Communication

An ORB knows whether an incoming request should be routed to local implementations or to another ORB running on a different machine. When a request reaches the ORB for which it is intended, the request is passed to an object adapter. The *Portable Object Adapter (POA)* forms a link between an object's implementation and its presence (reference) on the ORB.

It is the IDL compiler's job to turn IDL definitions into stub and skeleton files. The stubs and skeletons are all language and ORB dependent. In a case in which one is developing a heterogeneous application, the same IDL file may be used to generate the stubs and skeletons for each language and ORB implementation.

The stubs generated by the compiler will be used by the client processes to communicate with the server. A skeleton file is the companion of a stub. It is the skeleton's job to receive requests from the ORB, call the proper implementation, and return the results. [Ref 7]

3. CORBA Event Service

A standard CORBA request results in synchronous execution of an operation by an object. If the operation defines parameters or return values, data is communicated between the client and the server. In some scenarios, a more decoupled communication model between objects is required. For example, several documents are linked to a spreadsheet. The documents are interested in knowing when the values of certain cells have changed. When the value of one of the cell changes, the documents update their presentations accordingly. Furthermore, if a document is unavailable because of a failure, it is still interested in any changes to the cells and wants to be notified of those changes when it recovers. [Ref 6]

The CORBA event service decouples the communication between objects. The event service defines two roles for objects: the supplier role and the consumer role. Suppliers produce event data and consumers process event data. Event data are communicated between suppliers and consumers by issuing standard CORBA requests.

There are two approaches to initiating event communication in CORBA. They are called the *push model* and *pull model*. The push model allows a supplier of events to

initiate the transfer of the event data to consumers. The pull model allows a consumer of events to request the event data from a supplier. In the push model, the supplier is taking the initiative; in the pull model, the consumer is taking the initiative.

An event channel is an intervening object that allows multiple suppliers to communicate with multiple consumers asynchronously. An event channel is both a supplier and a consumer of events. Event channels are standard CORBA objects and communication with an event channel can be accomplished using standard CORBA requests. An event channel mediates the transfer of events between the suppliers and consumers as follows:

- The event channel allows consumers to register interest in events, and stores this registration information
- The channel accepts incoming events from suppliers.
- The channel forwards supplier-generated events to registered consumers.

Suppliers and consumers connect via the event channel and not directly to each other. From a supplier's perspective, the event channel appears as a single consumer; from a consumer's perspective, the event channel appears as a single supplier. In this way, the event channel decouples supplier and consumers.

Any number of suppliers can issue events to any number of consumers using a single event channel. There is no correlation between the number of suppliers and the number of consumers, and new suppliers or consumers can connect to more than one event channel.

a. The Push Model

In the push model a supplier generates events and actively passes them to a consumer. In this model a consumer passively waits for events to arrive.

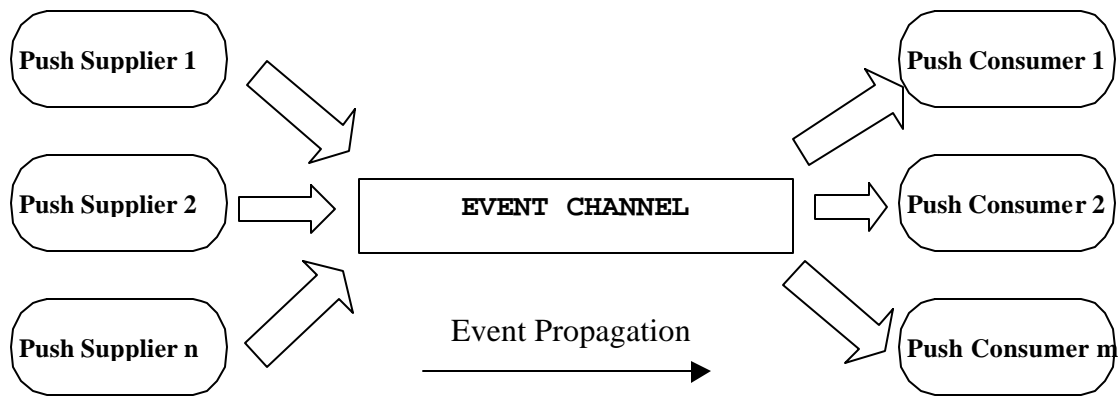


Figure 2.5 Push Model

In this architecture, a supplier initiates the transfer of an event by invoking an IDL operation on an object in the event channel. The event channel invokes a similar operation on an object in each consumer that has registered with the channel.

b. The Pull Model

In the pull model, a consumer actively requests that a supplier generate an event. In this model, the supplier waits for a pull requests to arrive. When a pull request arrives, event data is generated by the supplier and returned to the pulling consumer.

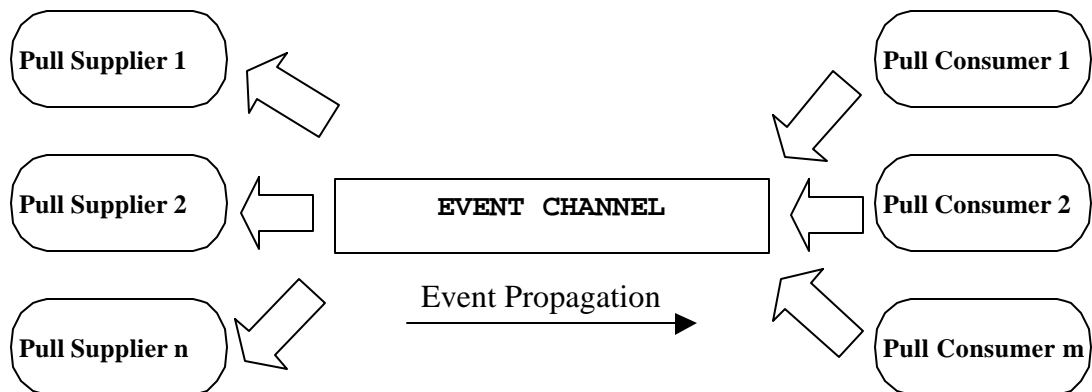


Figure 2.6 Pull Model

In this architecture, a consumer initiates the transfer of an event by invoking an IDL operation on an object in the event channel application. The event channel then invokes a similar operation on an object in each supplier. The event data is

returned from suppliers to the event channel and then from the event channel to the consumer, which initiated the transfer.

c. *Mixing the Push and Pull Model in a Single System*

Because suppliers and consumers are completely decoupled by an event channel, the *Push* and *Pull* models can be mixed in a single system. For example, suppliers may connect to an event channel using the *Push* model, while consumers connect using the *Pull* model, as shown in Figure 2.7.

In this case, both suppliers and consumers must participate in initiating event transfer. A supplier invokes an operation on an object in the event channel to transfer an event to the channel. A consumer then invokes another operation on an event channel object to transfer the event data from channel. Unlike the case in which consumers connect using the *Pull* model, the event channel takes no initiative in forwarding the event. The event channel stores events supplied by the push supplier until some pull consumer requests an event, or until a push consumer connects to the channel.

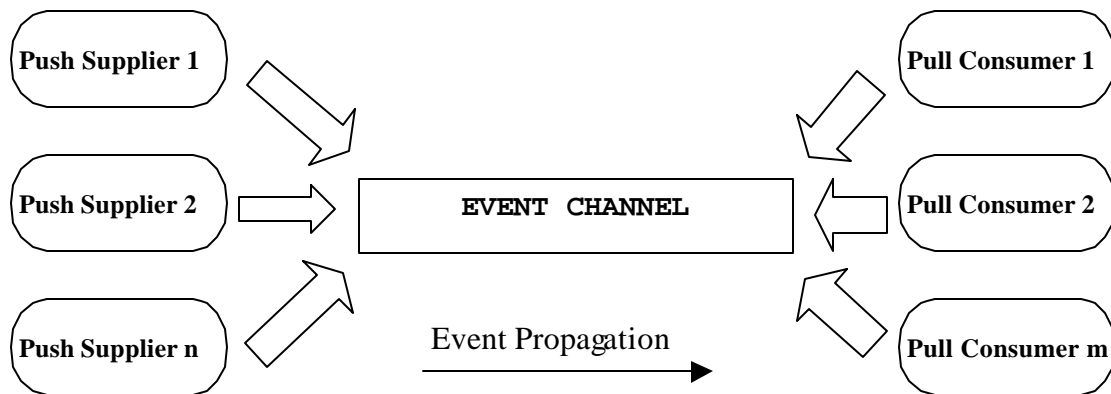


Figure 2.7 Mixed Model

d. *Types of Event Communication*

The CORBA Event Service maps an event to a successfully completed sequence of operation calls. The operation and the sequence of calls are clearly defined for both *Push* and *Pull* models, and data related to an event can be passed as operation

parameters or return values. Event communication can take one of two forms, *typed* or *untyped*.

(1) **Untyped Event Communication.** In untyped event communication, an event is propagated by a series of generic *push* or *pull* operation calls. The *push* operation takes a single parameter, which stores the event data. The event data parameter is of type *any*, which allows any IDL defined data type to be passed between suppliers and consumers. The *pull* operation has no parameters but transmits event data in its return value, which is also of type *any*. Clearly, in both cases, the supplier and consumer applications must agree on the contents of the *any* parameter and return value if this data is to be useful.

(2) **Typed Event Communication.** In typed event communication, a programmer defines application-specific IDL interfaces through which events are propagated. Rather than using *push* and *pull* operations and transmitting data using an *any* argument, a programmer defines an interface that suppliers and consumers use for the purpose of event communication. The operation defined on the interface may contain parameters defined in any suitable IDL data type. In the *Push* model, event communication is initiated simply by invoking operations defined on this interface. The *Pull* model is more complex because event communication is initiated by invoking operations on an interface that is specially constructed from application-specific interface that the programmer defines. Event communication is initiated by invoking operation on the constructed interface. [Ref 8]

4. CORBA Notification Service

CORBA Event Service falls short in several areas. There is no Quality of Service (QoS) support that can be utilized to set the certitude of notification transmission. Many times, event notifications need to be processed so consumers can meet requirements of reliability, priority, ordering, and timeliness. Also, there is no capability to filter events so a consumer will only get a subset of events. Moreover, existing Event Service implementations are usually filled with proprietary mechanisms. Once a set of events is

implemented within a system, it becomes impossible to integrate those events with another application that uses the same set of events but a different Event Service implementation. In order to make this service more transferable, progress has been made to achieve greater agreement on implementing some of these mechanisms that have become proprietary. There have been efforts to enhance the functionality and interfaces of the Event Service. The enhancements have resulted in a new CORBA service -- the Notification Service.

a. Filter

One of the most promising additions to the Notification Service is filtering. Filters allow consumers to subscribe to particular events by matching events to be delivered to them against constraint expressions. This allows the Notification Service to cut down on traffic and improve the scalability of CORBA event handling. Filters are CORBA objects that enable interfaces to add, remove, and modify constraints that match event message values. Constraints refer to variables that are bound to parts of the event notification message and are specified using event types and/or expressions written in a constraint language.

b. Quality of Service (QoS)

The Notification Service provides better support for QoS by defining standard interfaces that allow control over the notification delivery's Quality of Service characteristics. Name/value pairs are used to represent service characteristics at different levels of the protocol stack. Each characteristic has some impact on notification delivery. Therefore, at the channel, one may want to set a discard policy for determining which notifications are thrown away as resource limits are breached. Or one may want to set a maximum number of notifications that may be queued for a single consumer. At the message level, one may want to set the reliability of event delivery. This could be a "best effort" value, where no guarantees of delivery are made, or it could be set as "persistent," in which case the Notification Service stores the notification until the connection is re-established. Some of the Quality of Service properties defined by the specification include:

- **Discard policy:** Specifies policies for discarding events when the queues are full.
- **Earliest delivery time:** Specifies how long an event is to be held in the channel before it is delivered.
- **Expiration time:** 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.
- **Maximum events per consumer:** Provides an upper bound to the number of events the channel will queue on behalf of a given consumer. This QoS value helps to relieve the pressure that can grow within a channel because of misbehaving consumers that prevent some of the events from being consumed as scheduled. A huge number of notifications that are constantly being saved and queued for a single misbehaving consumer can use up a great deal of valuable system resources that would otherwise be available for well-behaved consumers.
- **Order policy:** Specifies the order in which notifications are buffered for delivery.
- **Priority:** Specifies the order in which events are delivered to consumers so that more important events take precedence.
- **Reliability:** Is linked to both event reliability and connection reliability to specify fault-tolerance properties. If these properties are in place, the Notification Service will reconnect to all its clients and deliver all non-expired events to its consumers after a crash.

c. Structured Events

With all this added information for each event, there is a need for a general event structure, in which a wide variety of event data can be stored. The structured event is a more strongly typed event message, known to the Notification

Service and its clients. This structure will allow more efficient processing of the event notification.

Structured events consist of a header and a body. The header contains two sections. The first has fixed information such as *domain_name*, *type_name* and *event_name*. The second section, which is the variable portion of the header, contains optional information about the event. This is constructed as a sequence of properties to hold QoS information related to the individual event notification. The separation of the header into two sections is designed so that a lightweight notification can be created.

The body of the structured event contains the actual event data, and is also divided into two parts: the filterable data, and the remaining, or payload, data. The filterable part is a sequence of properties. These properties contain the fields of the notification on which consumers are most likely to base filter decisions. All other data is contained in the remainder of the body. This data is of type CORBA *any*. Again, this design was created to streamline throughput. There is no reason that an event could not be filtered on data in the payload portion of the body, but performance would not be optimal.

Also complete contents of the notification could be contained within the optional header fields, leaving the body empty. This would allow a more streamlined event. The structured event was organized to make filtering as simple and efficient as possible while not forcing one to use it in a specific manner. [Ref 9]

5. Conclusion

The CORBA Event and Notification Services provide powerful capabilities to handle events in a distributed, heterogeneous environment. However, CORBA only specifies a set of conventions and protocols that must be followed by CORBA implementations - it is left to developers to implement these specifications. Further, the CORBA Event Model is an integrated part of CORBA, so that it is difficult to separate the event model from CORBA and make it a generic tool for other projects that do not use CORBA.

C. SAAM EVENT CHANNEL MODEL

The current SAAM channel model was designed by the SAAM research team, headed by Dr. Geoffrey Xie, and implemented by NPS graduates Dean Vrable and John Yarger in their thesis “The Server and Agent based Active network Management (SAAM) Architecture: Enabling Integrated Service”. The channel model provides event dispatching between SAAM components. It was designed to overcome some limitations of the Java event model.

One of the limitations is the order of instantiation. In the Java event model, the order of instantiation of objects is important. For instance, an event listener cannot register with an event source that has not been instantiated.

Another limitation has to do with event source replacement. If one event source is to take over notification of a certain type of event, listeners for that event must de-register with the old source and then register with the new source. This process involves a great deal of object passing and quickly becomes cumbersome.

1. SAAM Channel

A SAAM channel is a mechanism for delivering SAAM protocol events from one or more event sources to one or more event listeners (Figure 2.8). Similar to the CORBA event model, a SAAM channel decouples the sources and listeners of an event.

A SAAM channel is currently implemented as a Java class, and is thus allowed to have attributes that constitute the state of the channel and methods that affect the state. A channel contains a vector of *listeners* and a vector of *talkers*. Under this model, event listeners become listeners to a channel. They register with a channel as they would register with any event source. The key difference here is event sources now must register as talkers on the channel. To send an event notification, a registered talker merely calls the `talk` method of the desired channel. The channel then notifies all listeners on that channel.

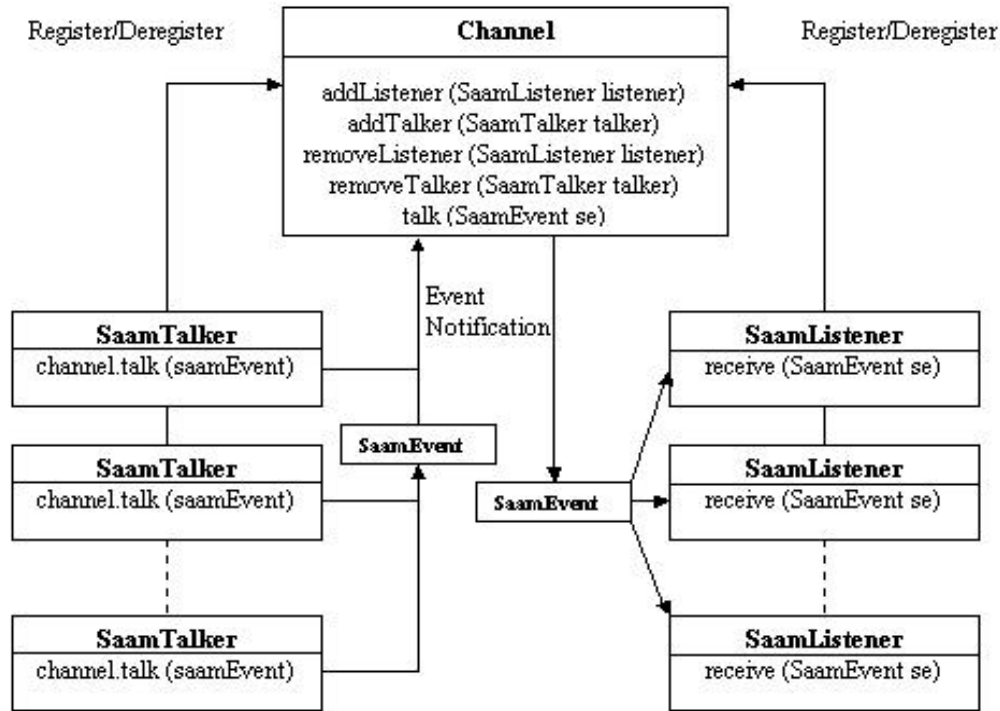


Figure 2.8 SAAM Channel Concept

In the channel registration process, the order in which objects register is not important. Listeners are allowed to register with a channel that has no talkers and vice versa. If a listener registers with a channel that has no talkers, the listener will not receive any events on that channel until a talker registers with the channel, and then talks to it. If a talker registers on a channel that has no listeners, the events sent by the talker will be dropped. [Ref 3]

In the SAAM prototype, a set of well-known channels with unique and static identification numbers is defined to provide the communication between SAAM components. Each channel serves a particular set of SAAM components that perform a well-defined task. For example, *channel 80004* is used to pass SAAM control messages to the *PacketFactory* component for processing. All components, which want to pass a SAAM control message to *PacketFactory*, may at any time register with this channel as a talker and pass their messages to the channel regardless of which instance of *PacketFactory* is listening. Therefore, associating channels with well-defined tasks and static identifiers facilitates loosely coupled communication between SAAM components.

The existing SAAM channel model offers several benefits over the standard Java Event Delegation model. However, there are some shortcomings in the current SAAM channel design and implementation. These deficiencies will be discussed in the next chapter.

THIS PAGE INTENTIONALLY LEFT BLANK

III. CHANNEL PERFORMANCE METRICS AND EVALUATION OF EXISTING SAAM CHANNEL MODEL

To motivate the new channel design, a set of performance metrics and qualitative objectives that help to characterize a channel model and implementation are defined in this chapter. The performance of the current SAAM channel model have been measured and qualified according to the defined metrics. The results are reported in this chapter. They raised several performance issues that have guided the design and implementation phase of the new channel model.

A. CHANNEL PERFORMANCE METRICS AND QUALITATIVE OBJECTITIVES

In this section, several channel performance metrics used in this thesis will be defined and explained.

1. Channel Throughput (R)

A channel's main task is to deliver events generated by registered senders (channel talkers) to registered receivers (channel listeners). The performance of this task can be measured by channel throughput, defined as the number of events delivered per second. An effective and efficient delivery mechanism is required to maximize the throughput.

The channel event delivery period is the time necessary for the channel to deliver an event to all registered channel listeners. It may vary from event to event. Also, when a listener uses a separate thread for event handling, delivery of an event to that listener means just passing the event to the listener object. It does not involve in the actual handling of the event. Therefore, channel throughput does not always measure how fast listeners process events. The maximum channel throughput is directly related to the event delivery time. The following explicitly defines the relationship between the maximum channel throughput and the minimum event delivery time:

R_{max} : *Maximum Channel Throughput*

T_d : *Minimum Channel Event Delivery Time*

$$R_{max} = \frac{1}{T_d}$$

2. Channel Work Rate (W)

Channel Work Rate is defined as the total number of events processed by the channel's listeners per second. Similar to the channel throughput, it is possible to improve channel work rate by using efficient event dispatching techniques. The key is to maximize the degree of parallelism in the handling of events by different listeners.

3. Channel Access Delay (D)

Channel Access Delay is defined as the latency between a channel talker's *talk request* and its actual access to the channel. A channel can only process one talk request at a time. It is assumed that a talk request will be blocked and put into a waiting queue if the channel is busy processing another talk request. The channel will serve talk requests in the queue following a specific algorithm (e.g., FIFO). The total waiting time a talk request spends in the queue is defined as the *Channel Access Delay of that Talk Request*. The average channel access delay of events from a talker is defined as the *Channel Access Delay of That Talker*. These delays are directly related to the total number of registered talkers and the talk frequency and period of each talker.

4. Event Talk Time (T_t)

Event Talk Time is the amount of time the channel spends in processing a talk request. This metric is important for event talkers as well as the channel itself. Excessive event talk time may decrease the performance of the talkers and the channel dramatically. This is because all talker threads will be blocked during each event talk period.

From the perspective of channel, access delay will increase and channel throughput will decrease as a result of long event talk time.

5. Thread Count (C)

Multi-threading a channel object may increase the performance of the channel with respect to the metrics defined above. However, it is not always a good idea to add more threads to the design of a channel. As systems differ with respect to the number of threads supported, for a channel design to be general purpose, it is important to make the channel scalable in terms of the thread count. Also, a large number of threads may incur a significant amount of resource overhead. Each `Thread` object consumes memory. In addition, each thread has two execution call stacks allocated for it by the Java Virtual Machine. One stack is used to manage Java method calls and local variables, while the other stack is used to keep track of native code.

A thread also consumes processor resources. There is inherent overhead in the scheduling of threads by the operating system. It happens when one thread's execution is suspended, its state stored, and another thread is given access to the processor so its execution may be resumed. This event is called a *context switch*. CPU cycles are required to do the task of context switching and can become significant if numerous threads are running. [Ref 10]

Another important issue, which should be considered when using threads, is synchronization of thread executions. The consequences of failing to properly synchronize threads which access shared resources are severe: data corruption and race conditions. These can cause programs to crash, produce incorrect results, or behave unpredictably. Even worse, these conditions are likely to occur only rarely and sporadically, making the problem hard to detect and reproduce. If the test environment differs substantially from the production environment, either in configuration or in load, these problems may not occur at all in the test environment, leading to the erroneous conclusion that the tested programs are free of major failures, when in fact the conditions triggering the failures simply have not yet been encountered.

Improper or excessive synchronization of threads can lead to other problems, such as poor performance and deadlock. While poor performance is certainly a less severe problem than data corruption, it can still be a serious problem. Writing good multithreaded programs requires walking a fine line, synchronizing enough to protect the

data from corruption, but not so much as to risk deadlock or impair program performance unnecessarily.

Synchronization comes at a cost, also. A synchronized block in the Java language is generally more expensive, in terms of execution time, than the critical section facilities offered by many platforms, which are usually implemented with an atomic "test and set bit" machine instruction. Even when a program contains only a single thread running on a single processor, a synchronized method call is still slower than an unsynchronized method call. If the synchronization actually requires contending for the lock, the performance penalty is substantially greater, as there will be several thread switches and system calls required. [Ref 11]

When adding additional threads to the design of the channel, the issues and costs must be considered carefully. At a minimum, synchronization of built-in channel threads should be transparent to a developer who uses the channel objects, and the developer should not be asked to synchronize an application specific (talker or listener) thread with a built-in channel thread in the code.

6. Adaptability of Channel

Adaptability measures the ease with which a channel model and implementation can be adapted to other projects and applications. A channel design with high adaptability meets software reusability and extendibility goals.

Reusability is the ability of software products to be reused, in whole or in part, for new applications. On the other hand, extendibility is the ease with which software products may be adapted to changes to the specification. [Ref 12]

An easily adaptable channel must be reusable for other projects and must be extendible according to the needs of developers and the requirements of applications.

7. Scalability of Channel

When the number of talkers and listeners registered with the channel grows or the individual load of channel participants increases, the channel should be able tolerate the

changes. The channel should be able to be configured by developers to solve scalability problems.

When applications and the number of necessary channels grow, another scalability issue arises. Developers should be able to add new channels to their application without unduly impacting performance and manageability.

8. Manageability of Channel

For small projects, the number of channels required will also be small and the manageability of the channels will not be a big concern.

On the other hand, large applications will require a relatively large number of channels and developers will have to spend a substantial amount of effort to manage these channels. Channels should be designed to minimize this effort.

Each channel object must have unique properties that differ from other channel objects. These properties will help developers to reach, retrieve, and trace the required channel objects in their code.

9. Functional Flexibility

The major goal of a channel is to deliver events generated by talkers to registered listeners in a fast, reliable manner.

However, a channel should be flexible enough to dynamically change the behavior of the event delivery mechanism. For example, event delivery order and service priorities of channel participants (channel talkers and listeners) can be configured or tuned in accordance with the needs of participant objects.

10. Ease of Use

A channel should provide an easily usable application-programming interface (API) for developers. The following discussions assume a nominal Java `channel` class that embodies this API.

a. Easy Creation of Channel

Developers should be able to create a channel with default properties by using the default constructor of `Channel`. The `Channel` class should provide overloaded constructors to create new channel objects with different attributes.

b. Meaningful Method Names

Choosing meaningful method names makes programs readable and helps avoid excessive use of comments. It also helps developers to explore the channel functionality readily. Method names must be related to tasks fulfilled by the methods.

c. Easy Configuration of Channel

The `Channel` class must provide class methods, which enable the channel to be configured or tuned to improve its performance and functionality.

d. Sufficient and Understandable Method Overriding

Method overriding eases the developers' task when programming with the `Channel` class. The `print` method of the Java `PrintStream` class is a good example for this metric.

e. Conflicts Between Methods

Semantics and functionality of the methods in the `Channel` class must be as clear as possible. Interactions between the methods must be well defined and well documented.

f. Conflicts Between Parameters in a Method

Method parameter names must be descriptive and must identify the objects expected. The name of a method and the type and order of its parameters generate the method signature. Method signatures must be designed to avoid confusion.

g. Documentation

The Channel class should be well documented so that developers can study, use, and modify it easily.

B. TEST OF SAAM CHANNEL

The existing SAAM channel was tested with respect to the defined performance metrics and qualitative objectives. In this section, the test results are reported.

1. Channel Throughput and Work Rate

Throughput and work rate of the existing SAAM channel was measured with different event handling periods for listeners. In the test scenario, three talkers and three listeners were registered with a SAAM channel (Figure 3.1). All channel talkers generated events based on *Poisson* distribution and the average event generation rate was 200 events per second. The event handling time of *Listener 1* and *Listener 2* was set to 0.5 milliseconds per event achieved by an idle *for-loop* in the listeners' *receiveEvent* methods. On the other hand, the event handling time of *Listener 3* was increased gradually by forcing the execution thread to sleep for an increasing amount of time in the listener's *receiveEvent* method.

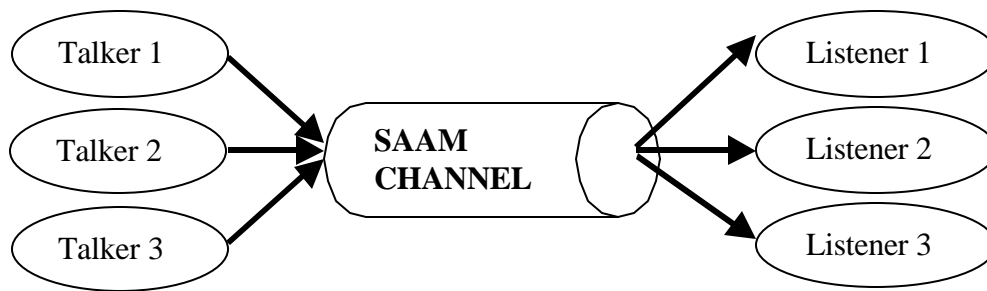


Figure 3.1 Test Bed For Studying Impact of Listener Event Handling Time on Channel Throughput

As shown in Figure 3.2, the channel throughput and work rate decreased severely when the event handling time of *Listener 3* was increased. These results show that the throughput and work rate of the existing SAAM channel is dependent on the total event delivery time, which depends on the number of registered channel listeners and the sum of their event handling times.

Event filtering can be used to multicast the events to interested listeners instead of broadcasting to all listeners. On the other hand, new threads can be created by the channel and dedicated to serve listeners that have a long event-handling time. Event filtering and dedicated event handling threads help reduce the total event delivery time and allow an event to be handled concurrently by different listeners. As a result, the channel throughput and work rate increases. The existing SAAM channel offers none of these capabilities and thus its throughput is not optimized.

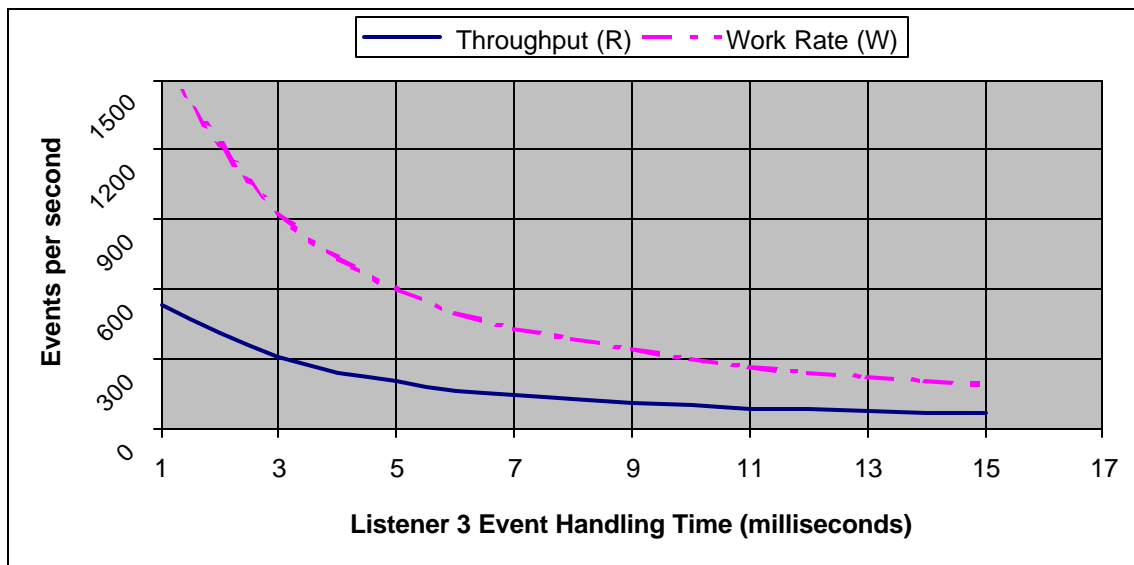


Figure 3.2 SAAM Channel Throughput and Work Rate versus Listener Event Handling Time

2. Channel Access Delay

Access delay of the existing SAAM channel was measured using a varying number of registered talkers and different event handling times for listeners. To measure the effect of the number of talkers on the channel access delay, three listeners with one millisecond event handling time (simulated by an idle *for-loop* in their *receiveEvent* method) were registered with the channel and the number of talkers was initially set to one and then increased steadily. Each talker generated 500 events per second, following a *Poisson* distribution. Figure 3.3 shows the test bed for the measurement.

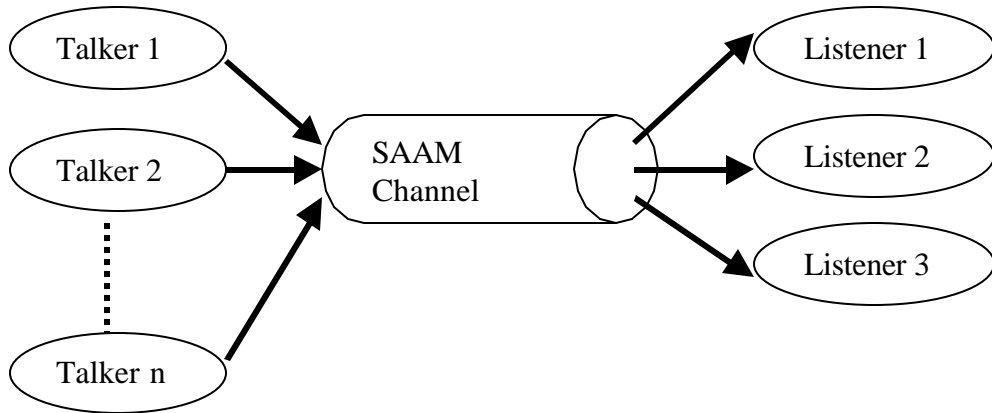


Figure 3.3 Test Bed for Measuring SAAM Channel Access Delay

The channel access delay for the case of a single talker was very small. This was an expected result because the channel was available to serve this talker all the time. The average channel access delay of talkers grew linearly when new talkers were added to channel. (Figure 3.4)

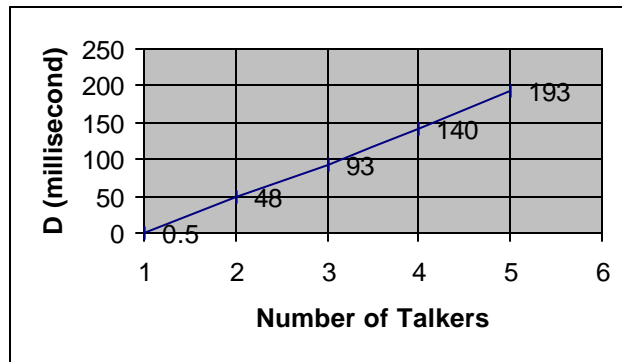


Figure 3.4 Average Access Delay versus Number of Channel Talkers

In the second scenario, the event handling time of *Listener 3* was increased gradually to see the effects of event handling time on the channel access delay. The rise in the total event delivery time caused the channel access delay to increase. (Figure 3.5)

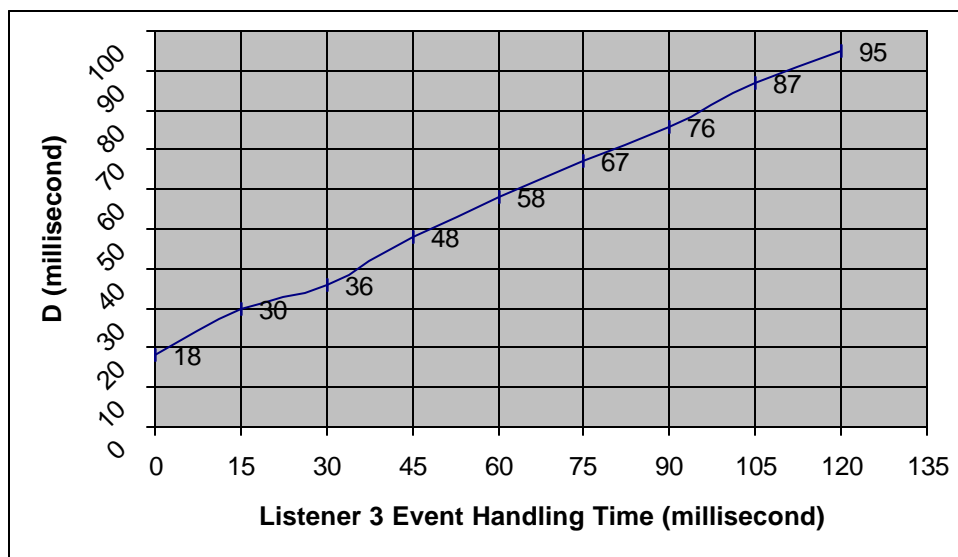


Figure 3.5 SAAM Channel Access Delay versus Event Handling Time

Therefore, the talker count and total event delivery time have a significant impact on the channel access delay of the existing SAAM channel. Specifically, the event handling times of channel listeners affect the channel access delay directly. This is because the existing SAAM channel has no buffering capability and it cannot accept another event until the current one is delivered to all listeners. This means that a channel is not accessible by talkers while it is dispatching an event. Consequently, the performance of the channel talkers is poor and unpredictable due to high and variable channel access delays.

3. Event Talk Time

As identified earlier, the existing SAAM channel has no event buffering capability and talkers cannot continue their job by sending their events to the channel. Talkers have to wait until the current event is delivered to all listeners. Actually, it is the thread of the current talker that executes channel dispatching and event handling code, except for those listeners, which are separately threaded.

The Event Talk Time of the SAAM channel includes the channel access delay and total event delivery time. Therefore, it should vary according to the number of talkers and

the event handling times of its listeners. A simple experiment was performed to verify this fact. The test bed depicted in Figure 3.3 was again used. Event Talk Time data was collected for an increasing number of talkers.

Figure 3.6 shows the results. Clearly, there is a linear rise in the event talk time when new talkers are added to the channel.

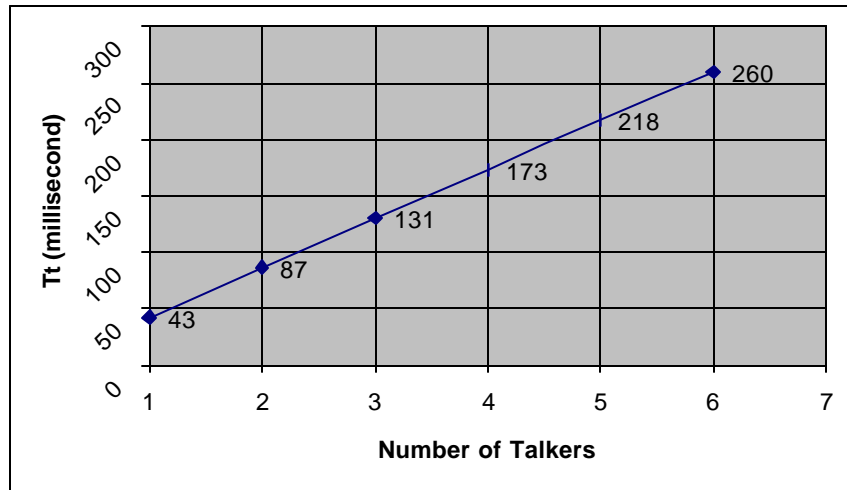


Figure 3.6 Event Talk Time Versus The Number of Talkers

More measurements of event talk time were collected using different event handling times. In this scenario, three talkers and one listener were registered with the SAAM channel. The event generation rates of the talkers and the event handling scheme of the listener are the same as the scenario used to study the channel access delay.

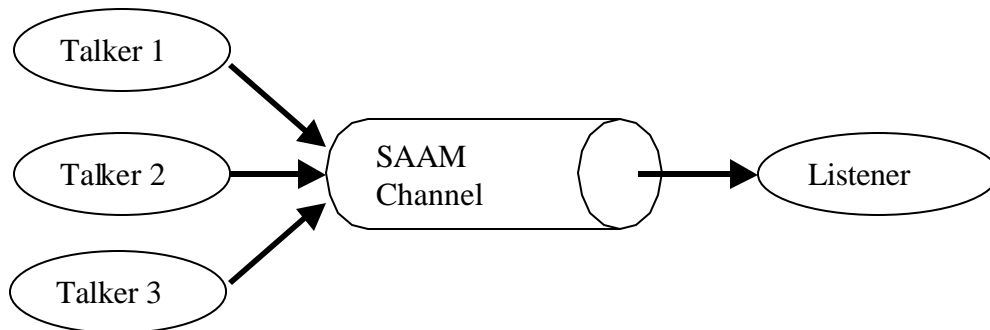


Figure 3.7 Test Bed for Measuring SAAM Channel Event Talk Time

The listener's handling time was increased steadily and the resulting event talk times recorded. Figure 3.8 shows the experiment results. It can be concluded that with the current SAAM channel design, event handling times affect event talk times directly.

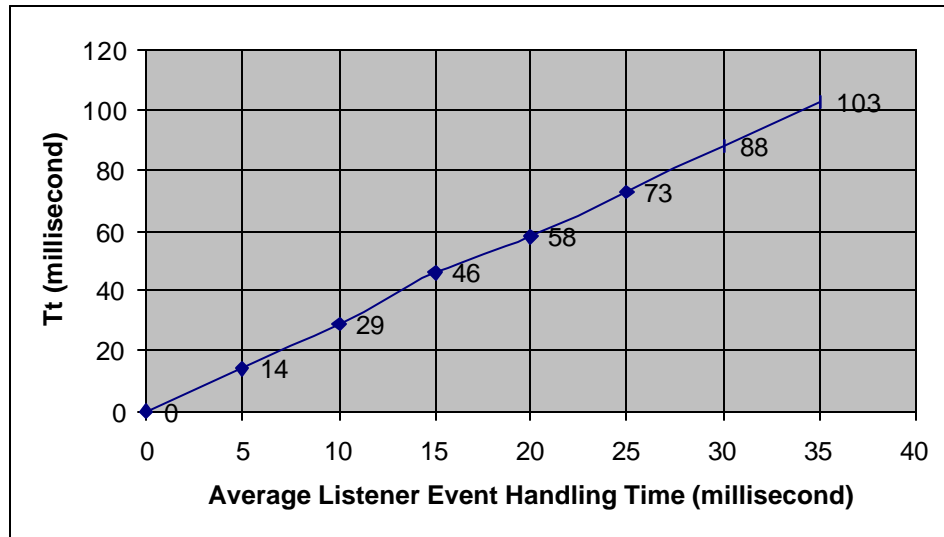


Figure 3.8 Event Talk Time Versus Event Handling Time

4. Thread Count

The existing SAAM channel does not require any thread to fulfill its functionality. This reduces its complexity. But, due to this threadless design, it cannot buffer events resulting in high and variable channel access delays and event talk times. The design may also cause some scalability problems, which will be discussed at the next section.

5. Manageability and Scalability of SAAM Channel

SAAM channels can be accessed by their channel identification numbers. An object wishing to manage these channels must create its own data structure, in order to keep records about existing channels and control the communication between them. In SAAM, the `ControlExecutive` class acts as a channel manager. SAAM talkers and listeners cannot register with or talk to channels directly. They interact with channels via `ControlExecutive`. `ControlExecutive` is implemented in an ad hoc manner. As a

result, the current SAAM channel design does not provide much support for channel management.

Scalability is one of the most critical drawbacks of the existing SAAM channel design. Adding new talkers and listeners to a channel increases the channel access delay and event talk time. The channel cannot compensate for these effects without performance penalties.

In addition, interaction between multiple SAAM channels should be considered carefully when the number of channels grows in an application. Consider the scenario depicted in Figure 3.9.

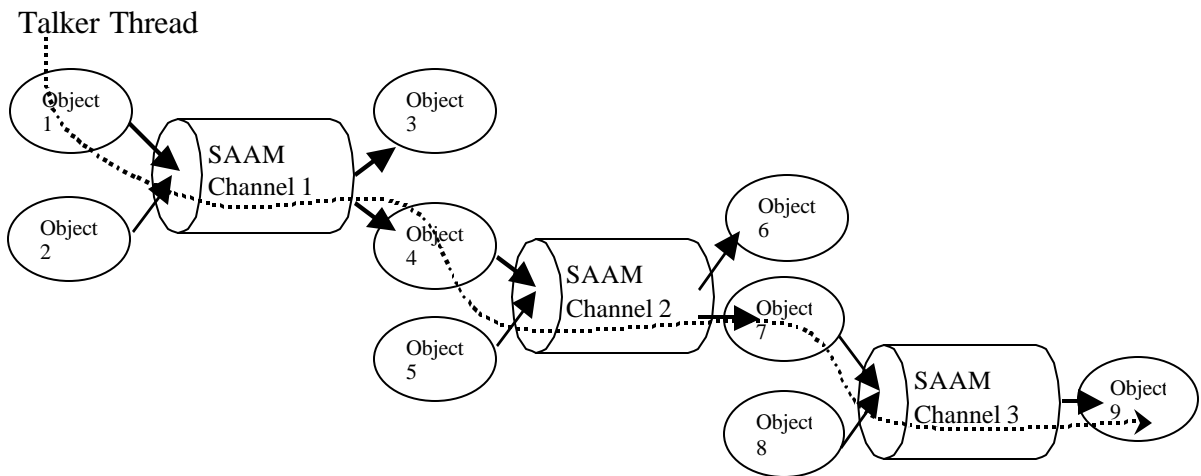


Figure 3.9 Event Concatenation in SAAM Channels

Object_1 sends an event to SAAM Channel_1. Object_4 talks to SAAM Channel_2 when it gets that event from Channel_1. Object_7 talks to SAAM Channel_3 when it gets the same event from Channel_2. This concatenation of channels blocks the thread of Object_1 until all three SAAM channels deliver the event to their listeners. At the same time, all SAAM channels in this scenario are inaccessible by other talkers until completion of Object_1's event delivery.

Another scalability issue is that the existing SAAM channel provides one-way communication. At least two channels are required to create a duplex communication between two objects. (Figure 3.10)

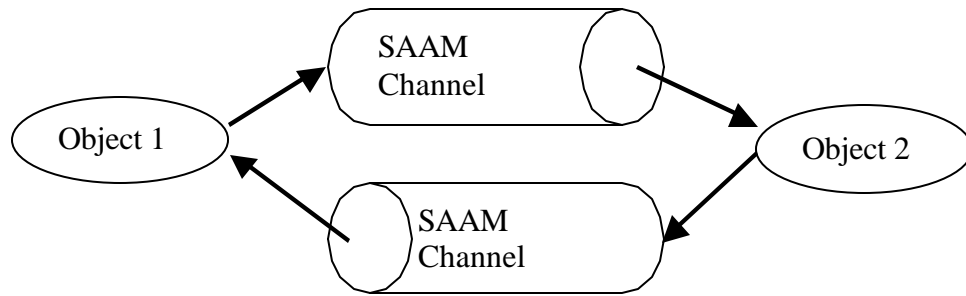


Figure 3.10 Two-way communication with SAAM channel

6. Adaptability and Functional Flexibility of SAAM Channel

The existing SAAM channel has been designed for communication between components of the SAAM prototype. Thus, it can handle only *SAAMEvents*. It is difficult to adapt the current SAAM channel model and implementation to other projects because it has been intended to meet the requirements of the SAAM prototype. It does not adapt to the varying needs of different applications.

Event dispatching of the existing SAAM channel is inflexible and cannot be changed. Events cannot be delivered to listeners in a specific order according to application needs. Listeners are served in their registration order and it is not possible to give precedence to listeners that have specific requirements for event handling.

In addition, the existing SAAM channels broadcast the current event to all registered listeners without regard to specific interests of the listeners. This increases the channel traffic unnecessarily and forces channel listeners to filter and drop uninterested events themselves.

7. Ease of Use of SAAM Channel

An instance of the SAAM channel can be created by providing an integer channel identification number and a channel participant, either a listener or talker. Later, other listeners and talkers can be added to the channel by using the *addListener* and *addTalker* methods. An additional constructor, which allows creation of a SAAM channel without any participant, would enhance the usability of the channel design.

IV. NEW CHANNEL MODEL AND CHANNEL PACKAGE

A. INTRODUCTION

The existing SAAM channel model offers several benefits over the standard Java Event Delegation model by allowing event delivery from one or more event sources to one or more event listeners and allowing event listeners to register with a channel that has no talkers.

However, there are some shortcomings in the current SAAM channel design and implementation. The model supports just one type of event class. It has no buffering capability. The thread of an event talker is blocked unnecessarily until the event is delivered to all listeners. The model does not give a programmer a choice of dispatching events to multiple listeners in a specific order. Consequently, the performance of the SAAM event model is not optimized.

As the main contribution of this thesis, a new channel model was designed to mitigate the shortcomings in the current SAAM channel implementation. The resulting channel package provides Java developers a generic, highly adaptable, and flexible communication mechanism between loosely coupled components.

B. FEATURES OF NEW CHANNEL DESIGN

The key features of the new channel design are summarized below. The implementation details of these features will be described in the next section where the Java channel package is introduced.

1. Generic Event Structure

The new channel supports multiple types of event objects per channel. Further, it uses a simple and extendible event structure.

2. Event Buffering

The new channel is capable of buffering events. This allows a channel to dispatch events to multiple listeners in a specific order and reduces the channel access delay and event talk time significantly.

3. Event and Channel Participant Priority

In the new channel design, a priority value is assigned to each event and each channel participant (channel talker and channel listener). Therefore, it is possible to give precedence to important events or serve particular listeners or talkers first.

4. Event Delivery Order

The new channel allows developers to choose different event dispatching schedules based on the priorities of events and channel participants. The dispatching can be customized based on application needs.

5. Event Filtering

The new channel design provides an event filtering mechanism, which allows channel listeners to subscribe to particular events. A channel will deliver to a listener only those events for which the listener is subscribed.

6. Self-Dispatching

The new channel design offers a self-dispatching mechanism for creating separate threads for channel listeners with large event handling times. This mechanism can be used to mitigate the negative effect on channel performance of long event handling times for specific listeners.

7. Duplex Communication

The new channel design supports two-way event communication between objects through a single channel. In other words, an object can be both a channel talker and a channel listener on the same channel, but for explicitly different events.

8. Concatenating Channels

The new channel design allows concatenation of channels. While a channel cannot generate events, it may be forwarded events from another channel. In this sense, a channel can be a listener on another channel.

C. CHANNEL PACKAGE

The new channel design was implemented in Java and all classes and interfaces were gathered under a Java utility package, called *channel*.

➤ Classes

- `Channel.java`
- `ChannelEvent.java`
- `ChannelListenerItem.java`
- `ChannelEventFilter.java`
- `FIFOScheduler.java`
- `PerTalker_RR_Scheduler.java`
- `PriorityScheduler.java`
- `ChannelManager.java`

➤ Interfaces

- `ChannelListener.java`
- `ChannelFilter.java`
- `ChannelScheduler.java`
- `ChannelAccessAuthority.java`

1. ChannelEvent Class

`ChannelEvent` class provides a data structure to encapsulate channel events. *Talker* and *event* data members of this class are instances of the `java.lang.Object` class, which is the base class of all Java objects. As such, it is the primitive class from

which all other classes are ultimately derived. Therefore, `ChannelEvent` class can encapsulate all objects as either an event or an event talker.

References to an event and its talker are required to create an instance of `ChannelEvent`. The value of the priority, for both the event and the talker, is an integer. If no value is provided, then a default value of zero is assigned. A higher integer value implies higher priority.

Another data member of `ChannelEvent` class is the timestamp. It is generated automatically, using the system's time in milliseconds.

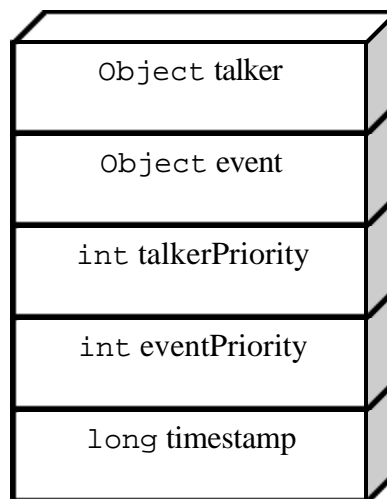


Figure 4.1 Channel Event Structure

2. ChannelListener Interface

All channel listeners must implement the `ChannelListener` interface to register with a channel and obtain events from that channel. `ChannelListener` contains a single method:

- `public void receiveEvent (ChannelEvent event)`

This method is called by the channel to deliver an event to the registered channel listeners. This method should be synchronized to guarantee that no more than one channel dispatcher thread is allowed when a channel listener registers with multiple

channels. The reason for this is that each channel creates its own event dispatcher thread and synchronization is critical to prevent the occurrence of hazards described earlier.

3. ChannelFilter Interface

Event filtering can be used by a channel to multicast its events to interested listeners instead of broadcasting to all listeners. Multicasting reduces the total delivery time for an event. Since the throughput is the reciprocal of the average event delivery time, the channel throughput will increase.

Without event filtering, channel listeners have to check all incoming events and choose the pertinent ones themselves. It can be guaranteed that channel listeners get only the events in which they are interested when event filtering is employed.

The `ChannelFilter` interface defines one simple method that must be implemented.

- `public boolean isAccepted (ChannelEvent event)`

Basically, if this method returns a value of true then this event passes the filter. Channel listeners can implement `ChannelFilter` themselves or use another object that implements this interface to take advantage of event filtering. Channel listeners can pass their filters to a channel when they register with the channel or later when filtering is needed.

Another benefit of event filtering is that channel listeners can control event flow between the channel and themselves dynamically. A listener can suspend or stop event acceptance from a channel without deregistering with channel. It is also possible to change the range and characteristics of accepted events to adapt to changes in system conditions and requirements.

The new channel design also allows concatenating event filters. A channel maintains the list of filters for a listener in a vector. New filters can be added to the end of the vector upon requests from the listener. The channel dispatches an event to that

listener only if this event is acceptable by all of the filters in the vector. Channel listeners can add a new filter or remove an existing filter by using the *addFilter* and *removeFilter* methods of `Channel`.

The `channel` package contains a built-in, ready-to-use, class named `ChannelEventFilter`, which implements the `ChannelFilter` interface and provides a filtering capability based on the class types of either the talker or the event, or simply the talker's name. `ChannelEventFilter` will be described in detail in Section 6 of this chapter.

4. **ChannelScheduler Interface**

A customizable channel scheduler allows a channel to buffer and dispatch events to multiple listeners in a specific order. A channel event can be categorized by either its talker, the event's type, or the respective priority of either the talker or the event itself. This categorization is used by the scheduler to determine the event delivery order.

The channel scheduler buffers and orders channel events according to its scheduling policy. The channel dispatches events to listeners by pulling them from the scheduler. The channel scheduler was implemented as an interface. The channel package includes three built-in channel schedulers; *FIFO Scheduler*, *Per Talker Round-Robin Scheduler*, and *Priority Scheduler*, all of which implement the `ChannelScheduler` interface. Developers may also implement and deploy their own schedulers to meet their specific requirements. This provides functional flexibility for the channel. The `ChannelScheduler` interface contains two methods:

- `public void push (ChannelEvent event)`
- `public ChannelEvent pull ()`

A channel forwards events to a scheduler by calling its `push` method and dispatches events by extracting them from the scheduler via the `pull` method.

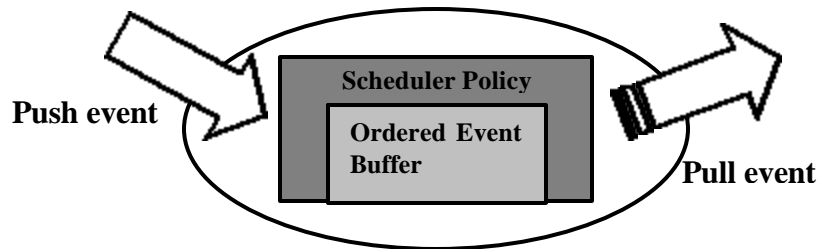


Figure 4.2 Channel Scheduler

5. ChannelListenerItem Class

ChannelListenerItem is the data structure that a channel uses to manage channel listeners. As depicted in Figure 4.3, each ChannelListenerItem object encapsulates a ChannelListener object, its filters, and its priority.

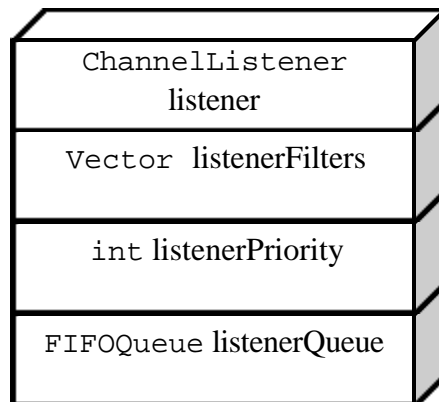


Figure 4.3 Channel Listener Item

The object also hosts an event buffer. The queue is used when a channel creates a dedicated thread for dispatching events for this listener. Without the option of having a separate dispatching thread, a listener with a long event handling time would delay the channel dispatcher thread and the other channel listeners unnecessarily. This self-

dispatching mechanism is implemented to eliminate a channel's dependence on the event handling periods of listeners.

After having registered with a channel, a listener may make a self-dispatch request by calling the *startListenerSelfDispatch* method of that channel. The channel relays this request to the channel listener item that it has created for the listener. The channel listener item allocates an event buffer and creates a thread for dispatching events from this event buffer to the listener. When the thread and the event buffer are ready, the self-dispatching flag of the channel listener is set to true. After it is set to true, the channel does not directly call the *receiveEvent* method of the listener for event delivery. Instead, the channel puts events into the event buffer of the corresponding listener item. If a channel listener has a variable event handling time in specific situations, self-dispatching can be suspended and resumed temporarily or can be stopped permanently.

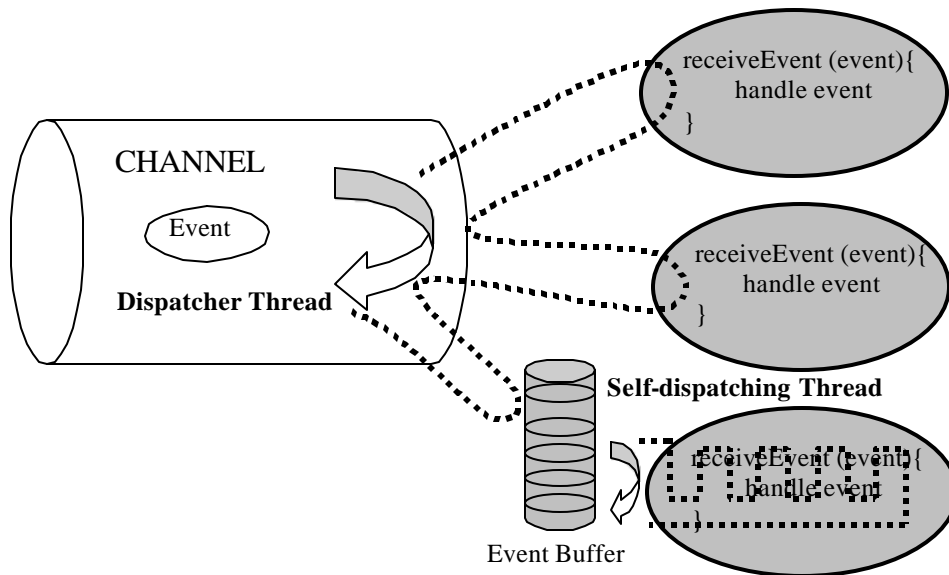


Figure 4.4 Self-dispatching

6. ChannelEventFilter Class

The channel package includes a built-in event filter class, *ChannelEventFilter*, which implements the *ChannelFilter* interface. Using this built-in filter, events may be filtered based on class types or talker name.

`ChannelEventFilter` class denies all events initially. In order for an event to be accepted, that is, to be passed by the filter, properties, which allow its acceptance, must be set by the listener. `ChannelEventFilter` provides four methods to set acceptance criteria:

- `addAcceptedTalkerName (String talkerName)`
- `addAcceptedTalkerClass (Class talkerType)`
- `addAcceptedEventClass (Class eventType)`
- `addAcceptedTalkerEventPair (Class talkerType,
Class eventType)`

The first method requires the name of the event talker as a parameter. The name of the talker is obtained by calling the `toString` method of the talker objects. All events whose talker names match this particular name will be accepted by the filter. The second and third methods are used to specify acceptable event class types and talker class types. The last method provides a more sophisticated filtering capability. A developer may specify acceptable events by their class types and their talker class types at the same time. The following methods are used to remove the accepted event properties, which were added previously.

- `removeAcceptedTalkerName (String talkerName)`
- `removeAcceptedTalkerClass (Class talkerType)`
- `removeAcceptedEventClass (Class eventType)`
- `removeAcceptedTalkerEventPair (Class talkerType,
Class eventType)`

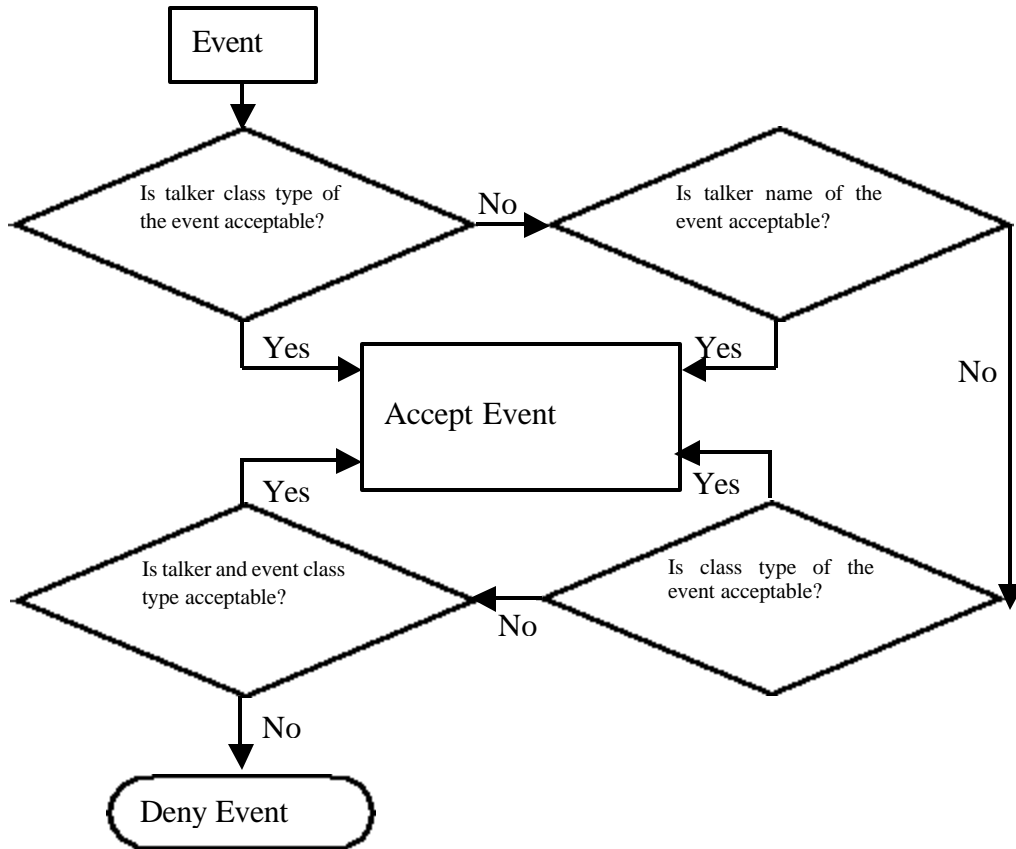


Figure 4.5 ChannelEventFilter Filtering Process

Figure 4.5 shows the filtering process for ChannelEventFilter class. There are also two useful methods: one to stop a filter from accepting events temporarily and the other to resume normal operation.

- suspendAccepting ()
- resumeAccepting ()

These methods do not change the event permission settings of the filter. All events are denied regardless of the filter settings when the *suspendAccepting* method is called. On the other hand, the *resumeAccepting* method returns the filter to its normal operation and events are accepted according to the existing filter settings.

7. Schedulers

As previously stated, the channel package contains three different built-in schedulers that implement `ChannelScheduler` interface: *FIFO Scheduler*, *Per Talker Round-Robin Scheduler*, and *Priority Scheduler*.

a. *FIFOScheduler Class*

The `FIFOScheduler` class simply buffers arriving events into a First In First Out (FIFO) queue. Events are served by their arriving orders, regardless of talker and event priorities. `FIFOScheduler` is the default scheduler for channels and will be installed when no scheduler is specified at the time of channel creation.

b. *PerTalker_RR_Scheduler Class*

The Per Talker Round-Robin Scheduler offers a fair service distribution between channel talkers. However, a priority queue is created for each talker. Arriving events are dispersed to separate queues based on their talker identity. The *pull* method extracts events from talker queues using a round-robin algorithm.

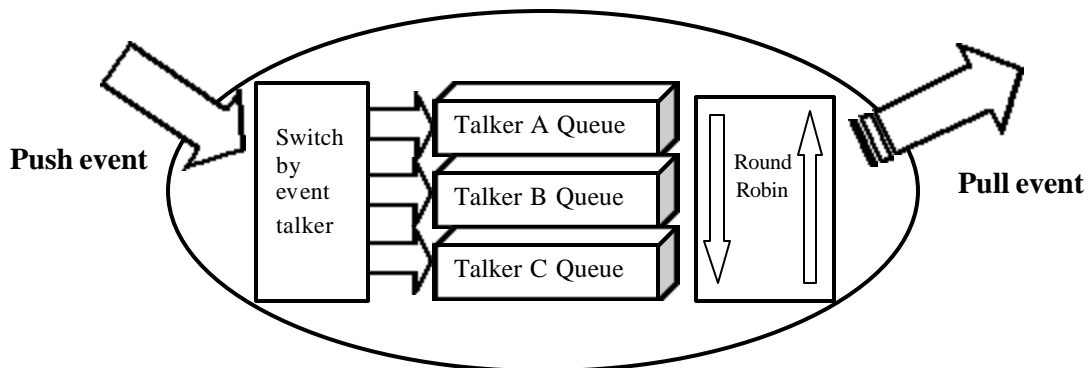


Figure 4.6 Per Talker Round-Robin Scheduler

c. *PriorityScheduler Class*

Events are buffered in a priority queue in accordance with their talker and event priorities. Events with the highest talker priority will be served first. If the talker priorities of two events are equal, then their event priorities are compared, and the one

with a higher priority will be served first. When both the event and talker priorities are equal, the first arriving event will be served first. This scheduler can be used to give event delivery precedence to particular event and talker types.

8. Channel Class

The Channel class is the core class that provides an event delivery service between channel talkers and channel listeners. Its main tasks include scheduling, filtering, and dispatching of events. Developers may customize how these tasks will be carried out for individual channels according to their application requirements.

a. Creating a Channel Object

There are two constructors to create an instance of Channel class:

- `public Channel (int channel_id)`
- `public Channel (int channel_id,
ChannelScheduler scheduler)`

The first one constructs a channel with a specified integer channel identification number. This constructor creates a channel with the default built-in FIFO scheduler. The second constructor can be used to create a channel with a specific scheduler. This scheduler can be chosen from one of the built-in schedulers or developed to meet particular requirements of the current application.

b. Adding and Removing Listeners to Channel

All objects that want to register with a channel as a listener must implement the `ChannelListener` interface. New listeners are added to a channel by invoking the channel's `addListener` method. This method is overloaded to allow the priority and filter objects of the listener to be specified as options.

- `addListener (ChannelListener listener)`
- `addListener (ChannelListener listener,
int priority)`
- `addListener (ChannelListener listener,
ChannelFilter filter)`
- `addListener (ChannelListener listener,
ChannelFilter filter, int priority)`

Existing registered listeners are removed from a channel by calling the following method:

- `removeListener (ChannelListener listener)`

c. Adding and Removing Talkers to Channel

Any object may register as a talker of a channel simply by invoking one of the channel's *addTalker* methods.

- `addTalker (Object talker)`
- `addTalker (Object talker, int talkerPriority)`

Talkers can be easily removed from a channel by using the following method:

- `removeTalker (Object talker)`

d. Talking on Channel

After registration with a channel, a talker calls one of the *talk* methods of the channel to send an event to the channel.

- `public void talk (Object talker, Object event)`
- `public void talk (Object talker, Object event,
int eventPriority)`
- `public void talk (ChannelEvent event)`

The first *talk* method is the most generic and the talker can use it to send any Java object to the channel as an event without restriction. The second would be used when the talker wants to assign a particular priority to the event. The third *talk* method makes it possible to deliver an instance of `ChannelEvent` directly to the channel. In the first two *talk* methods, the channel encapsulates the event object, the talker object, and the event priority (the default event priority is used if it was not specified) in a new `ChannelEvent`, while the third method requires a `ChannelEvent` as an input parameter. Each method pushes a `ChannelEvent` to the channel scheduler.

e. Event Dispatching

A dispatcher thread is created for a `Channel` object to pull events out of the channel scheduler and deliver them to each registered channel listener. Listeners are served in the order of their priorities. If two listeners have the same priority, the one that registered first will be served first (FIFO within priorities). The dispatcher thread continues to dispatch events continuously until the scheduler's event buffer is empty. Whenever the event buffer becomes empty, the dispatcher thread goes to sleep. It will be awakened upon arrival of a new event to the scheduler.

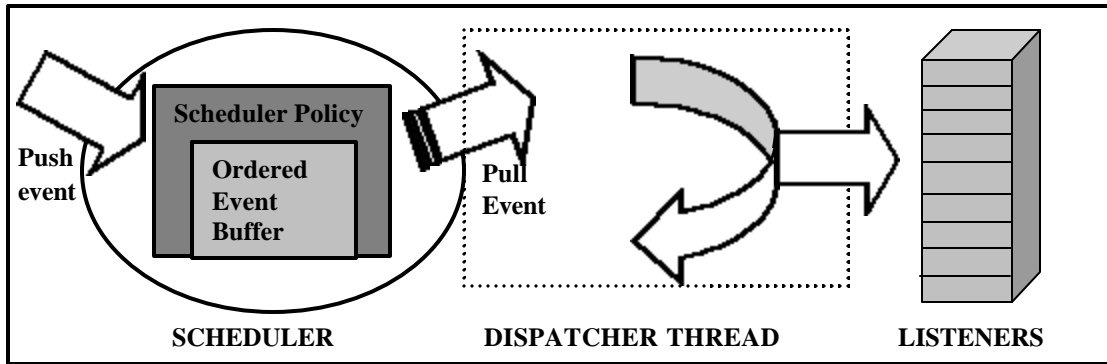


Figure 4.7 Channel Event Dispatching

f. Event Filtering

In the new channel design, a listener of a channel does not have to receive and handle all events that go through the channel. It can use event filters to identify specific types of events that it intends to receive. If one or more event filters are installed for a listener, the channel dispatcher always checks whether the current event is acceptable by those filters before delivering the event to the listener. If the event is not acceptable, the dispatcher will not deliver it to the listener. A listener is not required to specify a filter. Unless a listener specifies a filter, it will receive every event that goes through the channel. The use of a built-in filter with no specified acceptance criteria, as noted above, will prevent all events from being send to the listener. Figure 4.8 shows the event filtering process of a channel, which is repeated for every listener on a per event basis.

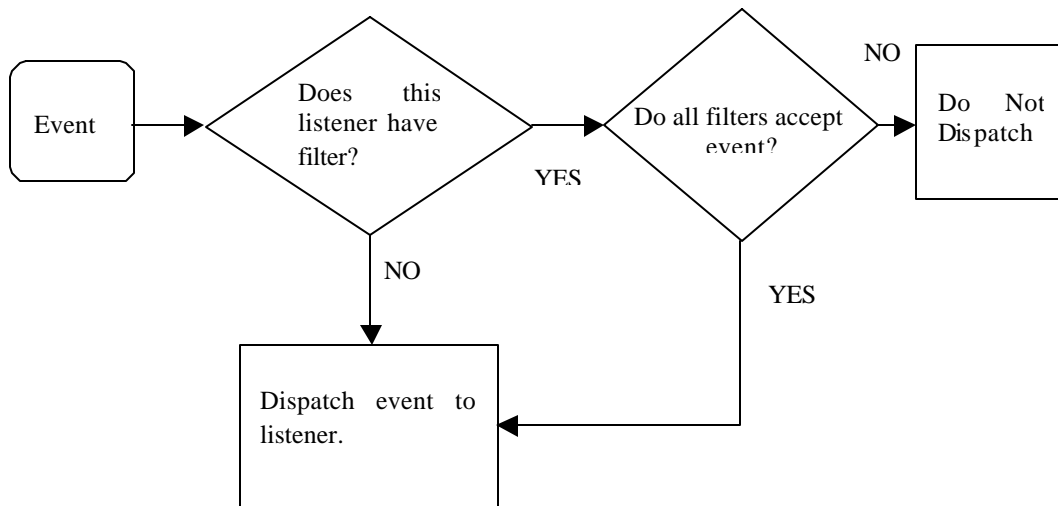


Figure 4.8 Channel Event Filtering Process

g. Listener Self-Dispatching

The ChannelListenerItem class supports self-dispatching, as discussed earlier. The Channel class provides the necessary API, i.e., public methods, to turn on and off self-dispatching for a listener. These methods are:

- startListenerSelfDispatch (ChannelListener listener)
- suspendListenerSelfDispatch (ChannelListener listener)
- resumeListenerSelfDispatch (ChannelListener listener)
- stopListenerSelfDispatch (ChannelListener listener)

A channel can start or resume self-dispatching for a listener without delay. However, suspending or stopping self-dispatching requires coordination between the channel dispatcher thread and the self-dispatching thread. The channel waits until the event buffer of the self-dispatching listener is empty to carry out a suspending or stopping request.

h. Duplex Communication

The new channel design supports two-way event communication between objects. In other words, an object can be a channel talker and channel listener on the same channel, but cannot receive its own events. Figure 4.9 illustrates the two-way communication between Object A and Object B.

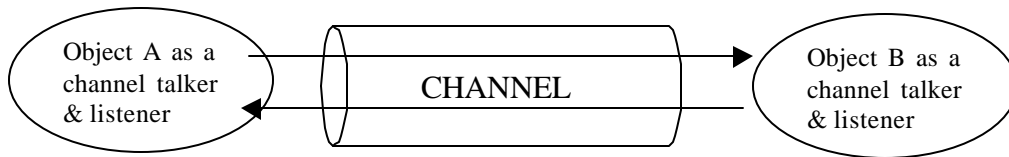


Figure 4.9 Two-way Event Communication

i. Concatenating Channels

The Channel class, itself, also implements the ChannelListener interface so that a channel can be a listener of another channel. This provides support for concatenating channels as shown in Figure 4.10. Only talkers can initiate events. A channel simply forwards events, submitted by a talker, to another channel as necessary, treating that receiving channel as a listener. The receiving channel must register as a listener using the source channel's *addListener* method.

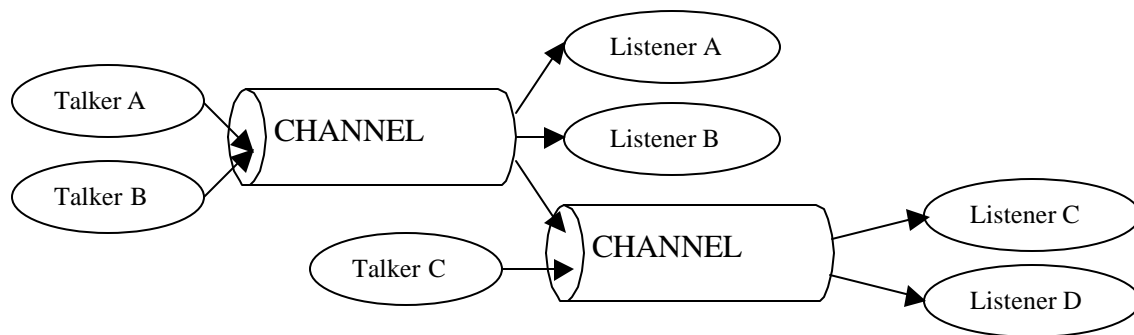


Figure 4.10 Concatenating Channels

9. ChannelAccessAuthority Interface

This interface provides a framework to control and organize all channels, and each channel's participants, of an application in a centralized manner. It defines five methods to be implemented by such a centralized channel controller.

- `boolean isTalkerAuthorized (Object talker, int channel_id)`
- `boolean isListenerAuthorized (ChannelListener listener int channel_id)`
- `int getTalkerPriority (String talkerName)`
- `int getListenerPriority (String listenerName)`
- `ChannelScheduler getSchedulerForChannel (int ch_Id)`

The first and second methods allow a channel controller component to specify permissions for talking and listening on all the channels. The controller may stipulate talker and listener priorities, the values of which may be retrieved via the third and fourth methods. The last method ensures that the proper channel scheduler is installed for a new channel. In this way, the channel authority can determine the event dispatching order for the channel in accordance with the requirements of the application and channel participants.

10. ChannelManager Class

The SAAM channel model completely decouples event sources (channel talkers) and event listeners (channel listeners). This decoupling raises an interesting question: "How can channel participants obtain a reference to the channel with which they want to interact?" This is not a big concern for small applications. Channel references can be passed to a participant object when the objects is created, or set later by calling the appropriate method of the channel participant of interest.

When the number of channels is large, it is necessary to create a common, easy-to-use, interface for use by all channel participants instead of passing channel references to participants individually. Channel talkers and listeners interact with this common interface to access their channels.

The `ChannelManager` class was developed to provide a common interface for channel participants in large applications. Basically, this class manages channels by keeping a table of existing channels, creating new channels when required, and overriding the `Channel` class methods with an extra argument to allow the desired channel to be identified simply by its `channel_id`. A `ChannelManager` can only control object access to channels if it contains a reference to a `ChannelAccessAuthority` object. A `ChannelManager` can only obtain this reference during instantiation and through the control object control access to channels and enforce predefined priorities for channel talkers and listeners. Thus two constructors are necessary:

- `public ChannelManager ()`
- `public ChannelManager (ChannelAccessAuthority authority)`

The first method constructs a channel manager without any access control authority while the second method creates one with an access control authority.

When an object wants to register with a channel it sends a request to the channel manager with the channel identification number included. If the channel manager contains no reference to a channel access control authority object, it proceeds to check whether a channel object with the given ID exists. If no such channel exists, the manager creates a new channel object with the specified channel ID and adds its information to the channel table. Once the target channel object is located, or created if necessary, the channel manager calls one of the channel's `addTalker` or `addListener` methods to complete the object registration.

However, if the channel manager contains a reference to a channel access control authority object, it first verifies that the requesting object is authorized access to the requested channel. This is accomplished by calling the control object's *isTalkerAuthorized* or *isListenerAuthorized* method. If the access is authorized, the channel manager will perform the same registration steps as described in the paragraph above. Otherwise, it will reject the request. Also, the channel manager will use the control object's *getTalkerPriority* or *getListenerPriority* method to set the priority of an authorized talker or listener. The required argument *talkerName* or *listenerName* is obtained by calling the requesting object's *toString* method. Finally, when creating a channel object, the channel manager will call the access control object's *getSchedulerForChannel* method to determine the appropriate channel scheduler for the new channel.

V. TEST AND RESULTS

Simple example applications have been built with the Java channel package. The performance of the new channel design was evaluated through experiments using these applications. The results are presented in this chapter. To demonstrate the suitability of the new channel design for large application development, the SAAM prototype has been updated to use the new channel package. The steps of this update are also described in this chapter.

A. EVALUATION OF NEW CHANNEL DESIGN

1. Channel Throughput and Work Rate

The new channel design supports event filtering and self-dispatching for channel listeners. When activated, both of these mechanisms may reduce the total delivery time for an event and allow events to be handled concurrently. As a result, the channel throughput and work rate should improve. The same experiment scenario used for measuring the SAAM Channel throughput and work rate was repeated for the new channel to verify this hypothesis. Figure 5.1 shows the test bed used in the experiment.

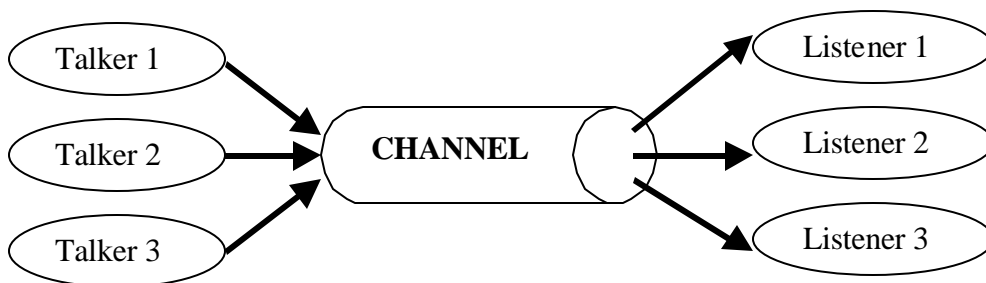


Figure 5.1 Test Bed For Studying Impact of Self-Dispatching on Channel Throughput and Work Rate

Self-dispatching was started for *Listener 3* when its event handling time reached 10 milliseconds. The channel throughput and work rate were monitored for the duration

of the experiment. The results are compared to those of the SAAM Channel in Figure 5.2a and Figure 5.2b.

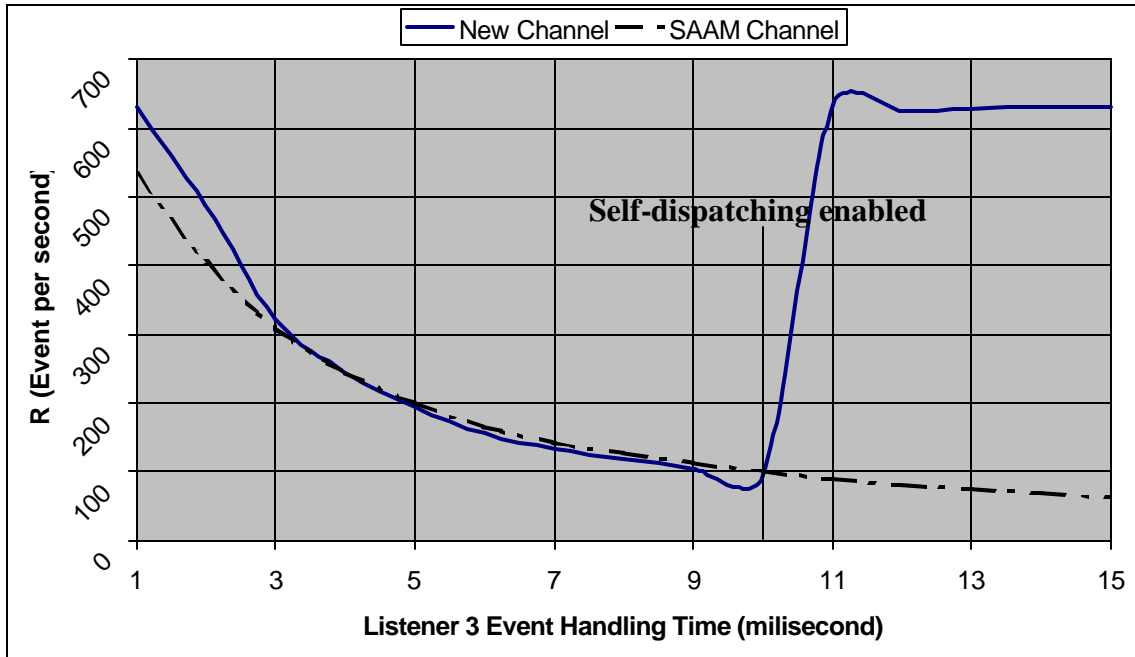


Figure 5.2a Effect of Self-Dispatching on Channel Throughput

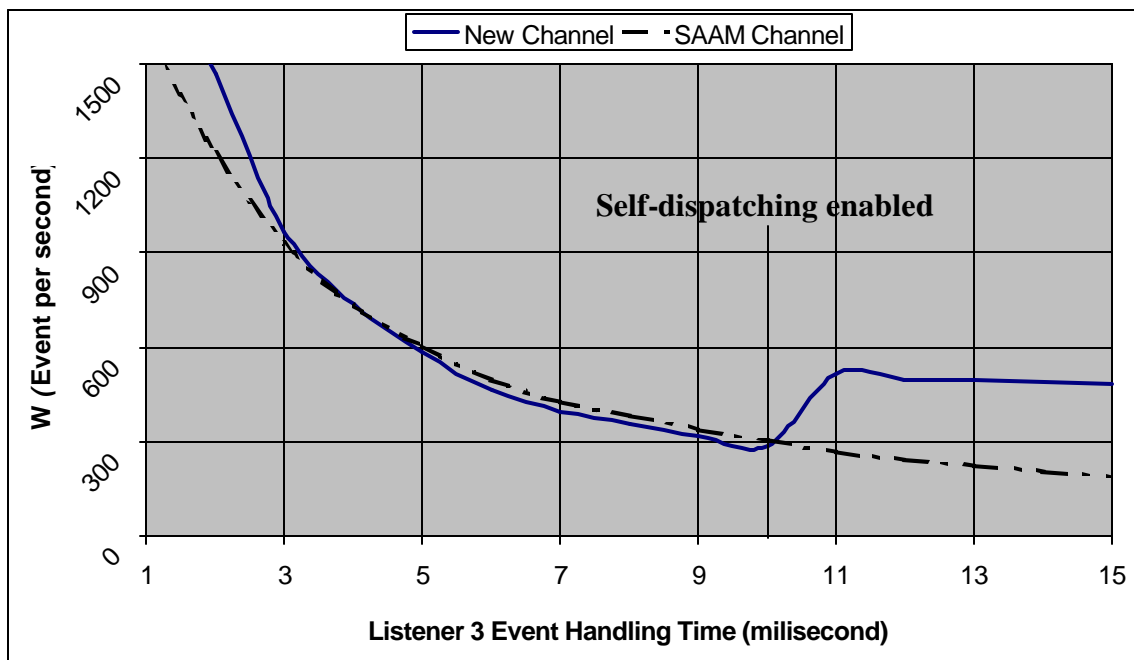


Figure 5.2b Effect of Self-Dispatching on Channel Work Rate

As anticipated, in the first phase of the experiment, the channel throughput and work rate decreased steadily as the event-handling time for *Listener 3* was increased without self-dispatching. As soon as self-dispatching was enabled for Listener 3, the average event delivery time decreased resulting in an increase to the channel throughput. The throughput quickly returned to the level observed before the event-handling time of Listener 3 was increased. Additionally, the channel work rate increased significantly and remained stable despite steadily increasing the event handling time of *Listener 3*.

As discussed in Chapter 3, a current SAAM channel's throughput and work rate depend on the event handling time of each of its listeners. The new channel allows a developer to mitigate this dependency by providing a self-dispatching mechanism.

2. Channel Access Delay and Event Talk Time

The new channel design reduces event talk times significantly by buffering events in the channel scheduler. Invoking the *talk* methods only incurs delay for pushing the event into the channel scheduler. Unlike in the existing SAAM channel model, a talker does not have to wait until its event is delivered to all listeners. The reduction of individual event talk time consequently causes channel access delays to decrease.

Another experiment was performed to evaluate the channel access delays and event talk times of the new channel design. The test bed shown in Figure 5.1 was again used. For the experiment, the number of talkers was increased gradually, and the corresponding average channel access delay and event talk time were measured.

The results are depicted in Figure 5.3 and Figure 5.4. The event talk time and channel access delay were a few microseconds in contrast to the millisecond range for the existing SAAM channel. These results show that the new channel design can handle a larger number of channel talkers without causing scalability problems.

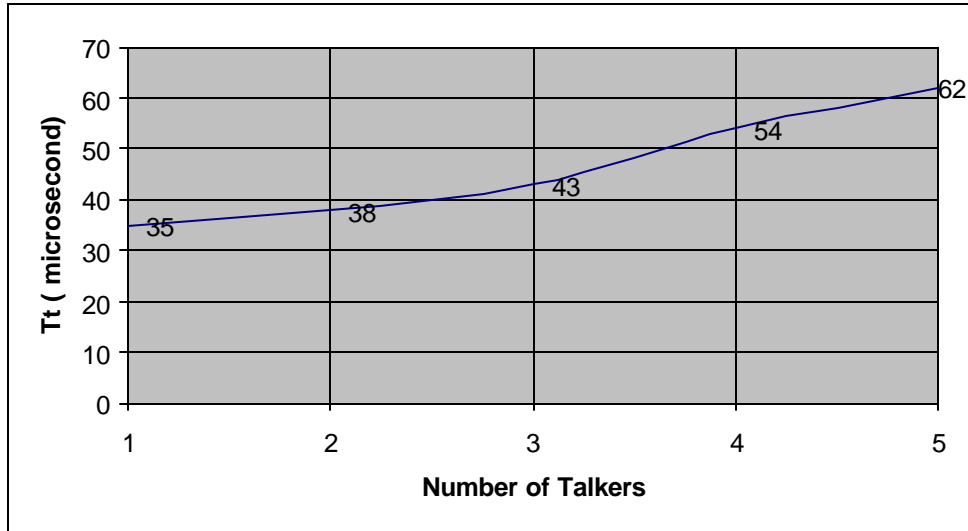


Figure 5.3 Event Talk Time versus Number of Talkers

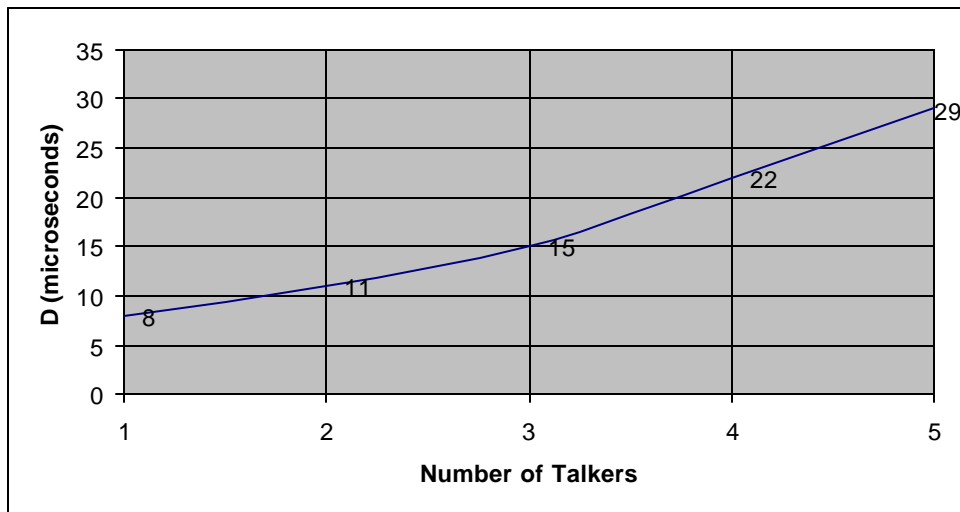


Figure 5.4 Channel Access Delay versus Number of Talkers

3. Scalability and Thread Count

In the new channel design, each channel employs one thread to perform its functionality. The thread dispatches events from the channel's scheduler to the channel's

listeners. This is necessary for event buffering, which is fundamental to minimizing the event talk time and channel access delay.

Since some systems can support only a limited number of threads, minimizing the channel (thread) count may be an important consideration. The new channel design addresses this scalability issue by providing support for event filtering and two-way communication. It does not restrict the number of participants for a channel and an object can be connected to both ends of a channel (being both a talker and listener) at the same time. As a result, it is feasible to reduce the number of required channels for an application, and thus the thread count, by increasing the number of participants on instantiated channels for that application.

4. Adaptability and Functional Flexibility of New Channel

The new channel design supports all event types (Java classes). It uses a simple and extensible event model instead of using a detailed, highly structured, and application-specific model. This generic event structure makes the channel package easily adaptable to either existing or new projects.

As discussed earlier, the new design partitions channel service into independent tasks, such as scheduling, filtering, and event dispatching. Developers may customize the functionality of each of these tasks to meet their specific needs. They may choose to use a particular scheduler for a channel. They may also use a specific set of event filters or activate self-dispatching for a listener. This functional flexibility makes the new channel configurable and tunable to meet different requirements for different projects.

5. Manageability

The new channel package offers a `ChannelAccessAuthority` interface and a `ChannelManager` class to help organize channels and control access of channel participants to these channels. They provide the necessary flexibility for managing the channels of a large application. A controller component based on them can offer a common interface for channel participants to request channel service, control the access to existing channels, and ensure proper priorities are assigned to channel participants.

6. Ease of Use

A new channel can easily be constructed with the default channel scheduler by providing a channel identification number.

Method names of the channel package were carefully chosen according to the tasks performed by the methods, in order to help developers understand the channel functionality. The documentation files of the channel package were created in HTML format by using the *javadoc* utility.

a. Easy Channel Management

In SAAM, a `ControlExecutive` object acts as a channel manager. SAAM talkers and listeners cannot register with or talk to channels directly. They can only interact with channels via the `ControlExecutive` for monitoring and delivery of all traffic on the existing channels. The `ControlExecutive` class overrides the methods of the SAAM Channel to achieve this purpose. However, this design adds extra cost and complexity to the `ControlExecutive` class.

One of the ideas behind the `ChannelManager` class in the new channel package is to provide a built-in and ready-to-use common interface for all channel participants when developers want to manage all channels and event traffic from a central component.

b. Easy Self-Dispatching

A listener with a long event handling time blocks the channel dispatcher thread and the other channel listeners. With the new channel design, developers can eliminate this problem by enabling self-dispatching for that listener.

The same situation also appears when programming with the Java Event Model. This is because all Java events are, by default, executed in a single thread, *the event dispatching thread*. Java creators documented this and left the creation and handling of new threads to programmers when there is a need to perform lengthy operations as resulting response to an event. In comparison, the self-dispatching

mechanism described herein is much more programmer friendly. All thread creation and synchronization issues are handled internally and re transparent to the developers.

B. INTEGRATION OF CHANNEL PACKAGE AND SAAM PROTOTYPE

In order to integrate the new channel design into the SAAM prototype, some modifications and additions were made to the existing SAAM code. In this chapter, these modifications and additions are described.

1. Removing Obsolete Classes and Interfaces

The new channel package was added into the `org.saamnet` directory. The following classes and interfaces that are part of the old channel implementation were removed from the existing SAAM code:

- `org.saamnet.saam.control.Channel` class
- `org.saamnet.saam.event.ChannelException` class
- `org.saamnet.saam.event.SaamListener` interface
- `org.saamnet.saam.event.SaamTalker` interface

The new channel event model implementation can encapsulate any Java object in an event. Therefore the current SAAM event and message structure were preserved. As depicted in Figure 5.5, SAAM events are now encapsulated in a generic `ChannelEvent` object. Preserving the existing event structure reduced the complexity of the integration effort.

All SAAM classes that implemented the `SaamListener` interface have been modified to implement the new `ChannelListener` interface.

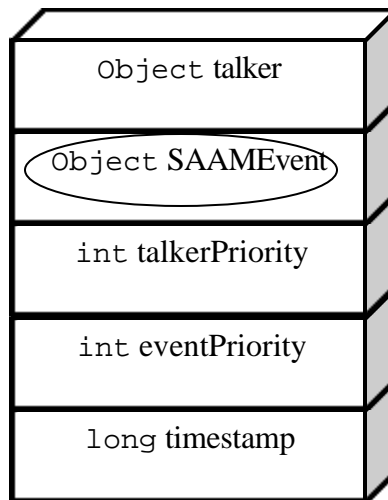


Figure 5.5 Encapsulating SAAM Events in ChannelEvent Structure

2. PermissionTableEntry Class

The `PermissionTableEntry` class was created to determine whether or not a channel participant has access to a given channel. A channel permission entry is created by providing the identification number of the channel and specifying the channel participants allowed to access this channel. The following method constructs an instance of `PermissionTableEntry` class:

- `PermissionTableEntry (int channel_id,
String [] allowedTalkers,
String [] allowedListeners)`

It contains three simple methods for accessing the entry attributes for a channel. These methods are `getValidChannel_Id`, `isValidTalker` and `isValidListener`. Additionally, its `getSaamPermissions` method returns a static permission table, as an array of permission table entries, containing the permissions for all channels, which are currently used in SAAM prototype, after the integration.

3. Changes to ControlExecutive Class

As mentioned earlier, the existing `ControlExecutive` class overrides methods of the SAAM Channel to control the monitoring and delivery of all traffic on the existing channels. This adds extra complexity to the `ControlExecutive` class.

For the integration, all channel management related methods were removed from the `ControlExecutive` class. An instance of the `ChannelManager` class was added to the `ControlExecutive` class as a new data member, providing a common interface for channel access. The `ControlExecutive` class was also modified to implement the `ChannelAccessAuthority` interface. When the channel manager object is instantiated it includes a reference to the `ControlExecutive` object as a channel access control authority object. A *getChannelManager* method was added to the `ControlExecutive` class to provide a reference to the embedded channel manager for other SAAM components.

All SAAM components now request channel service via the channel manager, with authorization performed by the `ControlExecutive` object.

4. Reducing the Number of Channels

The number of channels in the existing SAAM prototype depends on the number of interfaces and application agents in a node. When a new network interface is added to the node, a scheduler agent and a Network Interface Card (NIC) instance are created for the new interface automatically. In order to provide the event communication between these newly created components and the routing algorithm component, five new channels must be created. This scenario is illustrated in Figure 5.6 with each arrow representing a channel.

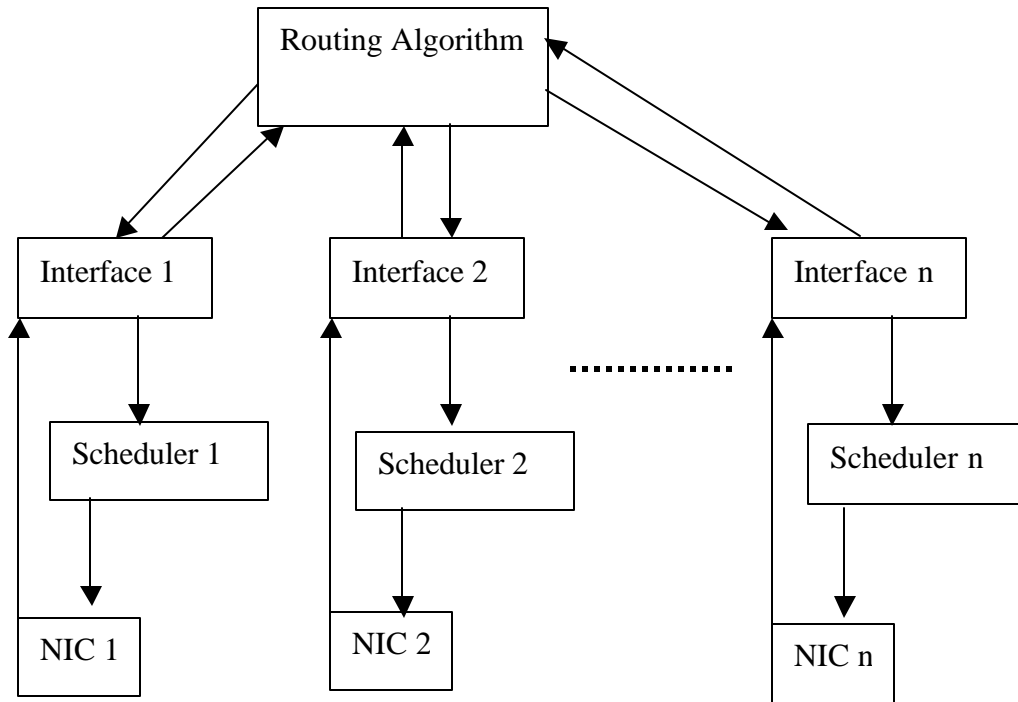


Figure 5.6 Necessary Channels for A New Interface on Existing SAAM Prototype

With integration of the new channel package to the SAAM prototype, a group of existing channels that are used for the same type of components were aggregated into one single channel by using the event filtering capability of the new channel design. A new channel is created only when new functionality (e.g., an application agent) is installed on a SAAM node. For example, only one channel is now used between all interfaces and their schedulers. Consequently, the number of channels for a given node no longer depends on the number of interfaces on that node.

Table 5.1 shows all the channels and their participants in the SAAM prototype after the integration. All talkers and listeners have the default priority of zero in each channel. To ensure fairness between registered channel participants a `PerTalkerRoundRobin` scheduler is used.

CHANNEL ID	CHANNEL TALKERS	CHANNEL LISTENERS	CHANNEL SCHEDULER
80010	Control Executive	Transport Interface	Priority
80020	Transport Interface All Interfaces	Routing Algorithm	Per Talker Round Robin
80030	Routing Algorithm	Transport Interface	Priority
80040	Routing Algorithm	All Interfaces	Priority
80050	All Interfaces	All Schedulers	Per Talker Round Robin
80060	All NICs Routing Algorithm	All Interfaces	Per Talker Round Robin
80070	All Schedulers	All NICs	Per Talker Round Robin
80080	All NICs	Translator	Per Talker Round Robin
80090	Translator TranslatorPortListener	All NICs	Priority
80100	Translator TranslatorPortListener Transport Interface	Packet Factory	Per Talker Round Robin

Table 5.1 New Channel Structure in SAAM Prototype

5. Event Priorities

SAAM fulfills its functionality by communicating control messages between SAAM nodes. Therefore, losses of control traffic should be minimized to increase the reliability of the SAAM network and to recover from component failures quickly. For example, the loss of a *flow-response message*, which is sent to a SAAM router by an active SAAM server, would cause network resources to be wasted.

In order to give precedence to SAAM control traffic, a two-level event priority scheme was defined for SAAM events in the `ControlExecutive` class:

- `public static final SAAMEVENT_PRIORITY_HIGH = 10;`
- `public static final SAAMEVENT_PRIORITY_LOW= 0;`

The `ControlExecutive` and `Translator` classes were modified to assign high priority to SAAM control traffic and low priority to data traffic. Other classes were also modified to propagate priorities of events within the SAAM protocol stack.

6. Additions to SAAM GUI

The `Channel` and `ChannelManager` classes record channel activities for debugging purposes. The existing SAAM GUI shows only the existing channel identification numbers and participants of these channels. The main GUI of a SAAM node was modified to show debugging records of channels and the channel manager as illustrated in Figure 5.7.

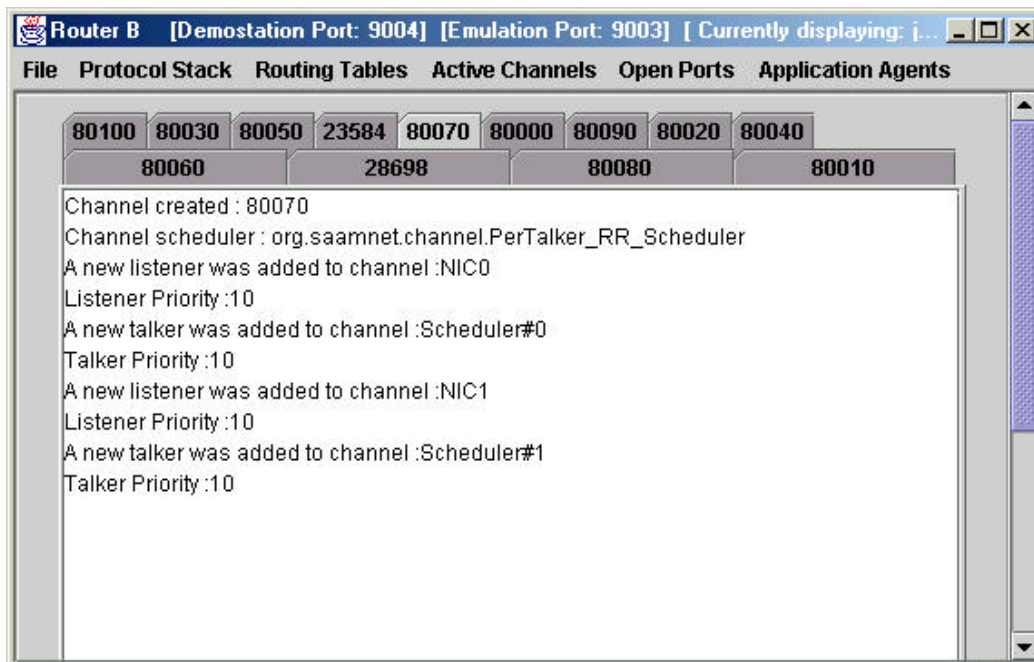


Figure 5.7 Snapshot of New Channel Debug Window

VI. CONCLUSION

In this thesis, a new, generic, highly adaptable, and flexible event channel was designed and implemented. The main product is a Java utility package, called “channel.” This package was designed to help Java application developers create or enhance large systems using an event-based programming approach. The new channel design has several demonstrated performance advantages over existing event channel implementations. The flexibility and adaptability of the channel package was also validated by a successful upgrade of the channel mechanism of the SAAM prototype system. The remaining sections describe several lessons learned from this thesis effort and identify potential future work required to enhance the channel package.

A. LESSONS LEARNED

1. Programming with Threads

In the implementation phase of this work, many thread issues arose. Although using multiple threads in a program offers important benefits, it also requires a good understanding of thread related concepts like synchronization, race conditions, and deadlock avoidance. This understanding took considerable time to acquire. Multi-thread programming appears, on the surface, to be easy with Java. However, it is difficult to get it right.

2. Integration with SAAM Prototype

Integration of the new channel package into the SAAM prototype turned out to be one of the more difficult parts of this thesis. This is because the SAAM prototype has more than two hundred Java classes and most of the SAAM components use event channels to communicate with each other. All of these components must be modified to work properly with the new channel design. It was necessary to understand the SAAM concepts and functionality completely in order to configure and test the new channels properly.

B. FUTURE WORK

1. Communication Between Distributed Applications

The current channel model provides an event communication mechanism between components of one application only. However, the channel model can be extended to support event communication between different, even distributed, applications, as well. To do so, it would be necessary to define an appropriate channel naming convention and develop a communication protocol to deliver events across process or network boundaries.

2. Automatic Sense for Self-Dispatching

For a channel, the event handling times of its listeners affect the channel throughput directly. Self-dispatching helps mitigate this problem. A channel listener may request its events to be dispatched by a separate thread by invoking the channel's *startListenerSelfDispatch* method.

Alternatively, an auto-sense mechanism can be embedded into the channel implementation to measure the event-handling time of each listener and automatically activate self-dispatching for a listener when specified criteria are met. In this manner, developers would no longer have to determine a priori which listeners require self-dispatching to optimize the throughput of a channel. This results in greater flexibility as it eliminates the need to hard-code when and where to initiate self-dispatching behavior on the part of specific listeners. Instead, the developers can focus on setting appropriate criteria for self-dispatching.

APPENDIX A. CHANNEL PACKAGE

```
package org.saamnet.channel;

import java.util.Vector;
import java.util.Hashtable;
import java.util.Iterator;
import java.util.Enumeration;
import java.awt.*;
import javax.swing.*;

/**
 * Channel is the core class that provides an event delivery service between
 * channel talkers and channel listeners. Its main tasks include scheduling,
 * filtering, and dispatching of events. Developers may customize how these
 * tasks will be carried out for individual channels according to their
 * application requirements.
 */

public class Channel implements ChannelListener,Runnable
{
    /**
     * Keeps the default priority for talkers
     */
    public static final int DEFAULT_TALKER_PRIORITY=0;

    /**
     * Keeps the default priority for listeners
     */
    public static final int DEFAULT_LISTENER_PRIORITY=0;

    /**
     * Keeps the identification number of this channel
     */
    private int channel_Id;

    /**
     * Keeps the number of delivered events
     */
    private int eventDelivered =0;

    /**
     * Keeps the number of accepted events
     */
    private int eventAccepted =0;
}
```

```

/**
 * Keeps the number of lost events
 */
private int eventLoss =0;

/**
 * Scheduler of this channel. Scheduler can be one of built-in schedulers
 * or developers can built new schedulers meeting their requirements by
 * implementing ChannelScheduler interface.
 * Scheduler accepts events from talkers and provides events to channel
 * dispatcher thread according to its scheduling algorithm.
 */
private ChannelScheduler scheduler;

/**
 * Keeps the talkers registered with this channel
 */
private Hashtable talkers;

/**
 * Keeps the listeners which are served by channel dispatcher thread.
 * Listeners will remove from this list when they are upgraded to
 * self-dispoatching.
 */
private Vector nonSelfDispatchListeners;

/**
 * Keeps the listeners which have their own dispatching threads.
 * Listeners will remove from this list when their self-dispatching
 * threads are suspended or stopped. Self-dispatching listeners get the
 * events earlier than non-selfdispatching listeners, regardless of
 * their priority
 */
private Vector selfDispatchListeners;

/**
 * Channel event dispatcher thread. This thread pulls the events from
 * channel scheduler and delivers them to listeners.
 *
 */
private Thread dispatcher;

/**
 * Keeps the status channel dispatcher thread;asleep or not
 */
private boolean dispatcherThreadAsleep = false;

/**
 * The synchronization lock for channel dispatcher
 */

```

```

private Object theDispatcherThreadLock= new Object();

/**
 * Keeps channel debugging statements
 */
private JTextArea debugText = new JTextArea();

/**
 * Constructs a channel with default built-in scheduler type.
 * Default scheduler type is FIFOSCHEDULER.
 * @param channel_Id Integer identification number for this channel
 */
public Channel(int channel_Id)
{
    this(channel_Id,new FIFOScheduler());
} // end of constructor

/**
 * Constructs a channel with custom scheduler. Developer can use
 * his/her own scheduler according to needs by implementing ChannelScheduler
 * interface.
 * @param channel_Id Integer identification number for this channel
 * @param scheduler Developer's custom scheduler
 */
public Channel(int channel_Id, ChannelScheduler sch )
{
    this.channel_Id = channel_Id;
    this.talkers = new Hashtable();
    this.nonSelfDispatchListeners = new Vector();
    this.selfDispatchListeners = new Vector();
    this.dispatcher = new Thread( this );
    this.scheduler = sch;
    String schedulerClass = (scheduler.getClass()).getName();
    if( schedulerClass.equals("org.saamnet.channel.PerTalker_RR_Scheduler") )
    {
        ((PerTalker_RR_Scheduler)scheduler).setParentChannel(this);
    }
    dispatcher.start();
    debugText.append("Channel created : "+ Integer.toString(channel_Id)+
        "\nChannel scheduler : "+ scheduler.getClass().getName()+"\n");
} // end of constructor

/**
 * Talkers use this method to push an event to channel without
 * an event priority.Channel encapsulates this event as an ChannelEvent
 * object with default priority (0).
 * @param talker Event owner
 * @param event Event object

```

```

*/
public void talk (Object talker, Object event)
{
    ChannelEvent newEvent = new ChannelEvent(talker, event);
    this.talk(newEvent);
} //end of talk() method

/**
 * Talker use this method to push an event to channel in form of
 * ChannelEvent object.
 * @param event An instance of ChannelEvent class
 */
synchronized public void talk (ChannelEvent event )
{
    if ( isRegisteredTalker(event.getTalker()) )
    {
        eventAccepted++;
        int talker_priority = ( (Integer) talkers.get(event.getTalker()) ).intValue();
        event.setTalkerPriority(talker_priority);
        boolean success = scheduler.push(event);
        if(! success) eventLoss++;
        synchronized (theDispatcherThreadLock)
        {
            if (dispatcherThreadAsleep)
            {
                theDispatcherThreadLock.notify();
            }
        }
    }
    else
    {
        String warning = "\nUnregistered Talker. Access Denied !" +
            "\nTalker: "+event.getTalker().toString()+
            "\nChannel Id: "+ Integer.toString(channel_Id);
        System.out.println (warning);
        debugText.append(warning+"\n");
    }
} //end of talk() method

/**
 * Talkers use this method to push an event to channel with its priority.
 * Channel encapsulates this event as a ChannelEvent object in behalf of
 * talkers
 * @param talker Event owner
 * @param event Event object
 * @param priority Event priority
 */
public void talk (Object talker, Object event, int priority)
{
    ChannelEvent newEvent = new ChannelEvent(talker, event, priority);
    this.talk(newEvent);
}

```

```

} //end of talk() method

/**
 * Channel also implements ChannelListener interface so a channel can
 * register with another channel as a listener. This makes channels
 * pipable to each others.
 * @param event Channel event
 */
public void receiveEvent( ChannelEvent event)
{
    this.talk(event);
} //end of receiveEvent() method

/**
 * Adds a new talker to this channel with priority
 * @param talker a channel talker
 * @param priority talker priority
 * @return Returns true if adding talker is successful else false.
 */
public boolean addTalker(Object talker, int priority)
{
    if( ! talkers.containsKey(talker) )
    {
        talkers.put(talker, new Integer (priority));
        String info = "A new talker was added to channel :"+talker.toString()+
            "\nTalker Priority : " + priority;
        debugText.append(info+"\n");
        return true;
    }
    else
    {
        String warning = "Adding talker to channel was failed."+
            "\nThis talker has already been added into channel"+
            "\nTalker : "+talker.toString();
        debugText.append(warning+"\n");
        return false;
    }
} //end of addTalker() method

/**
 * Adds a new talker to this channel with default priority
 * @param talker a channel talker
 * @return Returns true if adding the talker is successful else false.
 */
public boolean addTalker(Object talker)
{
    return this.addTalker(talker,DEFAULT_TALKER_PRIORITY);
} //end of addTalker() method

/**

```



```

* Adds a new listener to this channel in the form of ChannelListenerItem.
* ChannelListenerItem class encapsulates a channel listener with specified
* priority and a given filter object.
* @param listenerItem An instance of ChannelListenerItem
* @return Returns true if adding the listener is successful else false.
*/
public boolean addListener(ChannelListenerItem listenerItem)
{
    if( ! isExistingListener(listenerItem.getListener()) )
    {
        insertListener(nonSelfDispatchListeners,listenerItem);
        String info = "A new listener was added to channel :"+
            (listenerItem.getListener()).toString()+
            "\nListener Priority : " + listenerItem.getListenerPriority();
        debugText.append(info+"\n");
        return true;
    }
    else
    {
        String warning = "Adding listener to channel was failed."+
            "\nThis listener has already been added into channel"+
            "\nListener : "+(listenerItem.getListener()).toString();
        debugText.append(warning+"\n");
        return false;
    }
} //end of addListener() method

/**
* Adds a new listener to this channel without priority and filter
* @param listener A channel listener
* @return Returns true if adding the listener is successful else false.
*/
public boolean addListener(ChannelListener listener)
{
    ChannelListenerItem newListener = new ChannelListenerItem(listener);
    return this.addListener(newListener);
} //end of addListener() method

/**
* Adds a new listener to this channel with a specified priority
* @param listener A channel listener
* @param priority Listener's priority
* @return Returns true if adding the listener is successful else false.
*/
public boolean addListener(ChannelListener listener, int priority)
{
    ChannelListenerItem newListener = new ChannelListenerItem(listener,priority);
    return this.addListener(newListener);
} //end of addListener() method

/**

```

```

* Adds a new listener to this channel with a given filter object
* @param listener A channel listener
* @param filter Listener's filter
* @return Returns true if adding the listener is successful else false.
*/
public boolean addListener(ChannelListener listener, ChannelFilter filter)
{
    ChannelListenerItem newListener = new ChannelListenerItem(listener,filter);
    return this.addListener(newListener);
} //end of addListener() method

/**
* Adds a new listener to this channel with a specified priority and
* a given filter object
* @param listener A channel listener
* @param filter Listener's filter
* @param priority Listener priority
* @return Returns true if adding the listener is successful else false.
*/
public boolean addListener(ChannelListener listener, ChannelFilter filter, int priority)
{
    ChannelListenerItem newListener = new ChannelListenerItem(listener,filter,priority);
    return this.addListener(newListener);
} //end of addListener() method

/**
* Inserts a listener to specified listeners vector according to
* its priority. There are two possible vectors :
* nonSelfDispatchListeners and selfDispatchListeners
* @param listeners listeners vector which the listener will be inserted
* @param newListener A new listener
*/
private void insertListener( Vector listeners, ChannelListenerItem newListener )
{
    synchronized(listeners)
    {
        Iterator i = listeners.iterator();
        int index = 0;
        while(i.hasNext())
        {
            ChannelListenerItem currentListener = (ChannelListenerItem) i.next();
            if( currentListener.getListenerPriority() >= newListener.getListenerPriority() )
            {
                index ++;
            }
            else break;
        }
        listeners.insertElementAt(newListener,index);
    }
} //end of insertListener() method

```

```

/**
 * Removes a registered talker from this channel
 * @param talker Talker object to be removed
 * @return Returns true if removing of the talker is successful else false.
 */
public boolean removeTalker(Object talker)
{
    if(talkers.containsKey(talker) )
    {
        talkers.remove(talker);
        String info = "This talker was removed from channel"+talker.toString();
        debugText.append(info+"\n");
        return true;
    }
    else
    {
        String warning="There is no such registered talker on channel "+channel_Id+
            "\nTalker :"+ talker;
        System.out.println(warning);
        debugText.append(warning+"\n");
        return false;
    }
}
//end of removeTalker() method

```

```

/**
 * Removes a registered listener from this channel
 * @param listener listener object to be removed
 * @return Returns true if removing of the listener is successful else false.
 */
public boolean removeListener(ChannelListener listener)
{
    boolean isFound = false;
    if( isExistingListener(listener))
    {
        Iterator i = nonSelfDispatchListeners.iterator();
        while(i.hasNext())
        {
            ChannelListenerItem currentListener= (ChannelListenerItem) i.next();
            if( currentListener.getListener().equals(listener) )
            {
                i.remove();
                isFound = true;
                break;
            }
        }
    }

    if(! isFound)
    {
        Iterator j = selfDispatchListeners.iterator();
        while(j.hasNext())
        {

```



```

        item = currentListener;
        break;
    }
}
return item;
} //end of retrieveListenerItem() method

/**
 * Add the specified filter object to desired channel listener.
 * @param listener Channel listener
 * @param newFilter Filter object, which will be added to listener
 * @return Returns true if addition of filter is successful else false.
 */
public boolean addFilter(ChannelListener listener, ChannelFilter newFilter)
{
    ChannelListenerItem item = retrieveListenerItem(listener);
    if( item != null)
    {
        item.addFilter(newFilter);
        String info = "A new filter object is added to this listener :"+item.getListener();
        debugText.append(info+"\n");
        return true;
    }
    else
    {
        String warning = "There is no such listener to add a new filter"+
            "\nUnfound Listener :"+ item.getListener();
        System.out.println(warning);
        debugText.append(warning+"\n");
        return false;
    }
} //end of addFilter() method

/**
 * Remove the specified filter object from desired channel listener.
 * @param listener Channel listener
 * @param oldFilter Filter object, which will be removed from the listener
 * @return Returns true if removal of filter is successful else false.
 */
public boolean removeFilter(ChannelListener listener, ChannelFilter oldFilter)
{
    ChannelListenerItem item = retrieveListenerItem(listener);
    if( item != null)
    {
        item.removeFilter(oldFilter);
        String info = "The old filter object is removed from this listener :"+
            +item.getListener();
        debugText.append(info+"\n");
        return true;
    }
}

```

```

else
{
    String warning = "There is no such listener to remove a filter"+
        "\nUnfound Listener :"+ item.getListener();
    System.out.println(warning);
    debugText.append(warning+"\n");
    return false;
}
} //end of removeFilter() method

/**
 * Remove all filter objects from desired channel listener.
 * @param listener Channel listener
 * @return Returns true if removal of filters is successful else false.
 */
public boolean removeAllFilters(ChannelListener listener)
{
    ChannelListenerItem item = retrieveListenerItem(listener);
    if( item != null)
    {
        item.removeAllFilters();
        String info = "All filter objects is removed from this listener :"+
            +item.getListener();
        debugText.append(info+"\n");
        return true;
    }
    else
    {
        String warning = "There is no such listener to remove filters"+
            "\nUnfound Listener :"+ item.getListener();
        System.out.println(warning);
        debugText.append(warning+"\n");
        return false;
    }
} //end of removeAllFilters() method

/**
 * Gets the channel id.
 * @return channel id
 */
public int getChannel_ID()
{
    return channel_Id;
} //end of getChannel_ID() method

/**
 * Checks that talker is registered with this channel or not.
 * @param talker Talker object to be tested
 * return true if the talker is registered with this channel,
 * false otherwise
 */

```

```

public boolean isRegisteredTalker(Object talker)
{
    return talkers.containsKey(talker);
} //end of isRegisteredTalker() method

/**
 * Checks that listener is registered with this channel or not.
 * @param listener channel listener to be tested
 * @return Returns true if the listener is registered with this channel,
 * false otherwise
 */
public boolean isRegisteredListener(ChannelListener listener)
{
    return isExistingListener(listener);
} //end of isRegisteredListener() method

/**
 * Test if this channel has any registered talker
 * @return Returns true if this channel has any registered talker,
 * false otherwise
 */
public boolean hasTalkers()
{
    return (!talkers.isEmpty());
} //end of hasTalkers() method

/**
 * Test if this channel has any registered listener
 * @return Returns true if this channel has any registered listener,
 * false otherwise
 */
public boolean hasListeners()
{
    return (!(nonSelfDispatchListeners.isEmpty() && selfDispatchListeners.isEmpty()));
} //end of hasListeners() method

/**
 * Searches the listeners' vectors to figure out if this listener is an
 * existing listener
 * @param listener Channel listener to be tested
 * @return returns true if if this listener is an existing listener,
 * false otherwise
 */
private boolean isExistingListener(ChannelListener listener)
{
    boolean isFound = false;
    Iterator i = nonSelfDispatchListeners.iterator();
    while(i.hasNext())
    {
        ChannelListenerItem currentListener= (ChannelListenerItem) i.next();
        if( currentListener.getListener().equals(listener) )

```

```

    {
        isFound = true;
        break;
    }
}

Iterator j = selfDispatchListeners.iterator();
if(! isFound)
{
    while(j.hasNext())
    {
        ChannelListenerItem currentListener= (ChannelListenerItem) j.next();
        if( currentListener.getListener().equals(listener) )
        {
            isFound = true;
            break;
        }
    }
}
return isFound;
} //end of isExistingListener() method

/**
 * Starts self-dispatching for a listener.Channel starts self-dispatching
 * by using ChannelListenerItem object which encapsulates the listener.
 * ChannelListenerItem class is capable to create a dedicated
 * event dispatcher thread for this listener. After starting the
 * self-dispatching, channel will not call listener's receiveEvent method
 * anymore.Channel will put the coming events into the listener's event queue,
 * self-dispatching thread will deliver the events in this queue.
 * This listener will also transferred from nonSelfDispatchListeners list
 * to selfDispatchListeners list.
 *
 * @param listener Channel listener to be upgraded to self dispatching
 * @return Returns true if starting self-dispatching is successful, else false.
 */
public boolean startListenerSelfDispatch(ChannelListener listener)
{
    if(isExistingListener(listener))
    {
        Iterator i = nonSelfDispatchListeners.iterator();
        while(i.hasNext())
        {
            ChannelListenerItem currentListener= (ChannelListenerItem) i.next();
            if( currentListener.getListener().equals(listener) )
            {
                i.remove();
                currentListener.startSelfDispatching();
                insertListener(selfDispatchListeners,currentListener);
                break;
            }
        }
    }
}

```



```

    }
    String info = " Self-dispatching is started for this listener :"+listener;
    debugText.append(info+"\n");
    return true;
}
else
{
    String warning = "There is no such registered listener to"+
        "start self-dispatching"+ "\nUnfound Listener :"+listener;
    System.out.println(warning);
    debugText.append(warning+"\n");
    return false;
}
}
} //end of startListenerSelfDispatch() method

/**
 * Suspends self-dispatching for a listener. Channel suspends the self-
 * dispatching gracefully. It waits until the listener's event queue is
 * empty. This guarantees that listener will get events in order and
 * listener will not get an event from channel dispatcher and self
 * dispatcher at the same time.
 * To fulfill this requirements, channel sets suspend-request flag true
 * and then channel dispatcher checks appropriate situation to take over
 * the dispatching.
 * This process will not end the self-dispatching thread. Self-dispatching
 * thread waits for resuming or stopping request.
 *
 * @param listener Channel listener to be downgraded to channel dispatching
 * by suspending
 * @return Returns true if suspending self-dispatching is succesfull, else false.
 */
public boolean suspendListenerSelfDispatch(ChannelListener listener)
{
    if(isExistingListener(listener))
    {
        Iterator i = selfDispatchListeners.iterator();
        while(i.hasNext())
        {
            ChannelListenerItem currentListener= (ChannelListenerItem) i.next();
            if( currentListener.getListener().equals(listener) )
            {
                currentListener.setSuspendRequest(true);
                break;
            }
        }
        String info = " Self-dispatching is suspended for this listener :"+listener;
        debugText.append(info+"\n");
        return true;
    }
    else
    {

```

```

String warning = "There is no such registered listener to"+
    "suspend self-dispatching"+"\\nUnfound Listener :"+listener;
System.out.println(warning);
debugText.append(warning+"\\n");
return false;
}
} //end of suspendListenerSelfDispatch() method

/**
 * Resumes a suspended self-dispatcher.
 *
 * @param listener Channel listener to be upgraded to self-dispatching
 * by resuming its existing suspended self-dispatcher
 * @return Returns true if resuming self-dispatching is succesfull, else false.
 */
public boolean resumeListenerSelfDispatch(ChannelListener listener)
{
    if(isExistingListener(listener))
    {
        Iterator i = nonSelfDispatchListeners.iterator();
        while(i.hasNext())
        {
            ChannelListenerItem currentListener= (ChannelListenerItem) i.next();
            if( currentListener.getListener().equals(listener) )
            {
                currentListener.setResumeRequest(true);
                break;
            }
        }
        String info = " Self-dispatching is resumed for this listener :"+listener;
        debugText.append(info+"\\n");
        return true;
    }
    else
    {
        String warning = "There is no such registered listener to"+
            "resume self-dispatching"+"\\nUnfound Listener :"+listener;
        System.out.println(warning);
        debugText.append(warning+"\\n");
        return false;
    }
} //end of resumeListenerSelfDispatch() method

/**
 * Stops self-dispatching for a listener. Channel stops the self-
 * dispatching gracefully. It waits until the listener's event queue is
 * empty. This guarantees that listener will get events in order and
 * listener will not get an event from channel dispatcher and self
 * dispatcher at the same time.
 * To fulfill this requirements, channel sets stop-request flag true
 * and then channel dispatcher checks appropriate situation to take over

```

```

* the dispatching.
* This process will end the self-dispatching thread.
*
* @param listener Channel listener to be downgraded to channel dispatching
* by stopping
* @return Returns true if stopping self-dispatching is succesfull, else false.
*/
public boolean stopListenerSelfDispatch(ChannelListener listener)
{
    if(isExistingListener(listener))
    {
        Iterator i = selfDispatchListeners.iterator();
        while(i.hasNext())
        {
            ChannelListenerItem currentListener= (ChannelListenerItem) i.next();
            if( currentListener.getListener().equals(listener) )
            {
                currentListener.setStopRequest(true);
            }
        }
        String info = " Self-dispatching is stopped for this listener :"+listener;
        debugText.append(info+"\n");
        return true;
    }
    else
    {
        String warning = "There is no such registered listener to"+
            "stop self-dispatching"+ "\nUnfound Listener :"+listener;
        System.out.println(warning);
        debugText.append(warning+"\n");
        return false;
    }
}
} //end of stopListenerSelfDispatch() method

/**
* Gets registered channel talkers
* @return channel talkers
*/
public Vector getTalkers()
{
    Vector talkerVector = new Vector();
    Enumeration e = talkers.keys();
    while(e.hasMoreElements())
    {
        talkerVector.add(e.nextElement());
    }
    return talkerVector;
}
} //end of getTalkers() method

/**
* Gets registered channel listeners

```

```

* @return channel listeners
*/
public Vector getListeners()
{
    Vector copy = (Vector)selfDispatchListeners.clone();
    copy.addAll(nonSelfDispatchListeners);
    return copy;
} //end of getListeners() method

/**
* Returns channel id as a string
* @return channel id
*/
public String toString()
{
    return "Channel"+ channel_Id ;
} //end of toString() method

/**
* This method is used by channel dispatcher thread for pulling events
* from channel scheduler and then these events are sent to "dispatch" method
* for delivering to listeners.
* If scheduler is empty, channel dispatcher will sleep until arriving
* of the scheduler's new event notification.
*
*/
public void run(){
    if(Thread.currentThread() != dispatcher)
    {
        throw new RuntimeException( " This is a self-running object, this"+
            " method can be called by only object itself");
    }

    while(true)
    {
        ChannelEvent current = scheduler.pull();
        if(current == null)
        {
            try
            {
                synchronized ( theDispatcherThreadLock )
                {
                    dispatcherThreadAsleep = true;
                    theDispatcherThreadLock.wait();
                    dispatcherThreadAsleep = false;
                }
            }
        }
        catch( InterruptedException ie)
        {
            System.out.println(ie.toString());
        }
    }
}

```

```

    }
    else
    {
        dispatch(current);
        eventDelivered++;
    }
}
} //end of run() method

```

```
/**
```

```

 * This method is complimentary part of "run" method. It delivers
 * events to listeners. If listener is a non-selfdispatching listener
 * it will call the "receiveEvent" method of listener. If listener is
 * a self-dispatching listener, it will put the event into listener's
 * event queue for self-dispatching instead of calling "reciveEvent"
 * method.
 * Self-dispatching listeners are served earlier than non-selfdispatching
 * listeners, regardless of their priority. Because self-dispatching
 * listeners need much more time to handle events and we want self-
 * dispatching listeners to begin event handling as soon as possible.
 * This approach doesn't effect the non-selfdispatching listeners much
 * because channel dispatcher spends time for just putting events into self-
 * dispatching listeners' event queue.
 * This method also deals with the stop/suspend/resume requests for self
 * dispatching. It provides a graceful transition between self-dispatching
 * non-selfdispatching.

```

```
*
```

```
* @param item event to be dispatch
```

```
*/
```

```
void dispatch(ChannelEvent item)
```

```

{
    //Serve self-dispatching listeners and deal with control request(suspend & stop)
    Vector selfCopy = (Vector) selfDispatchListeners.clone();
    Iterator j = selfCopy.iterator();
    while(j.hasNext())
    {
        ChannelListenerItem currentListener = (ChannelListenerItem)j.next();
        if( currentListener.isAcceptedEvent(item)
            && !( currentListener.getListener() ).equals(item.getTalker() ) )
        {
            if(currentListener.isSelfDispatchQueueEmpty() &&
                (currentListener.hasSuspendRequest()||
                 currentListener.hasStopRequest()))
            {
                if(currentListener.hasSuspendRequest())
                {
                    currentListener.setSuspendRequest(false);
                }
                if(currentListener.hasStopRequest())
                {
                    currentListener.setStopRequest(false);
                }
            }
        }
    }
}

```



```

/**
 * Gets the debug text
 * @return Debug text area
 */
public JTextArea getDebugText()
{
    return debugText;
} //end of getDebugText() method

/**
 * Gets the number of accepted events
 * @return Returns the number of accepted events
 */
public int getAcceptedEventCount()
{
    return eventAccepted;
} //end of getAcceptedEventCount() method

/**
 * Gets the number of accepted events
 * @return Returns the number of delivered events
 */
public int getDeliveredEventCount()
{
    return eventDelivered;
} //end of getDeliveredEventCount() method

/**
 * Gets the number of accepted events
 * @return Returns the number of accepted events
 */
public int getLostEventCount()
{
    return eventLoss;
} //end of getLostEventCount() method

} //end of Channel class

```

```
package org.saamnet.channel;
```

```
/*
```

```
*ChannelEvent class provides a data structure to encapsulate channel events.
```

```
*Talker and event data member of this class are instances of java.lang.Object class.
```

```
*Therefore,ChannelEvent class can encapsulate all objects as an event and event talker.
```

```
*An event and its talker are required to create an instance of ChannelEvent.
```

```
*Event and talker priorities are optional and default event and talker priorities are
```

```
*assigned when they are not specified. Timestamp is generated automatically according
```

```
*to the system's time in milliseconds.
```

```
*/
```

```
public class ChannelEvent
```

```
{
```

```
    /*
```

```
     * Keeps the talker of event
```

```
    */
```

```
    private Object talker;
```

```
    /*
```

```
     * Keeps the class type of talker
```

```
    */
```

```
    private Class talkerClass;
```

```
    /*
```

```
     * Keeps priority of the event
```

```
    */
```

```
    private int eventPriority;
```

```
    /*
```

```
     * Keeps priority of the talker
```

```
    */
```

```
    private int talkerPriority;
```

```
    /*
```

```
     * Keeps the event that is encapsulated by this class
```

```
    */
```

```
    private Object event;
```

```
    /*
```

```
     * Keeps the class type of event
```

```
    */
```

```
    private Class eventClass;
```

```
    /*
```

```
     * Keeps instantiation time of event
```

```
    */
```

```
    private long timestamp;
```



```

/*
 * Used to direct an event to specified listener when needed
 */
private int listenerIndex;

/*
 * Keeps the default priority for events
 */
public static final int DEFAULT_EVENT_PRIORITY=0;

/*
 * Constructs a channel event with default talker and event priority
 * @param talker The talker of event
 * @param event Event object to be encapsulated
 */
public ChannelEvent (Object talker, Object event)
{
    this.talker = talker;
    this.talkerClass = talker.getClass();
    this.eventPriority = this.DEFAULT_EVENT_PRIORITY;
    this.talkerPriority = Channel.DEFAULT_TALKER_PRIORITY;
    this.event = event;
    this.eventClass = event.getClass();
    this.timestamp = System.currentTimeMillis();

}

} // end of constructor

/*
 * Constructs a channel event with a specified event priority
 * @param talker The talker of event
 * @param event Event object to be encapsulated
 * @param ev_priority Priority of event
 */
public ChannelEvent (Object talker, Object event, int eventPriority)
{
    this.talker=talker;
    this.talkerClass=talker.getClass();
    this.eventPriority=eventPriority;
    this.talkerPriority = Channel.DEFAULT_TALKER_PRIORITY;
    this.event=event;
    this.eventClass=event.getClass();
    this.timestamp = System.currentTimeMillis();

}

} // end of constructor

/*
 * Sets priority of the event.
 * @param priority Event priority

```

```

*/
public void setEventPriority (int priority)
{
    this.eventPriority=priority;
} //end of setEventPriority() method

/*
 * Sets priority of the talker. Channel sets talker priorities according to
 * talker registration.
 * @param priority Talker priority
 */
public void setTalkerPriority (int priority)
{
    this.talkerPriority=priority;
} //end of setTalkerPriority() method

/*
 * Sets listener's index for this event.
 * @param index listener index
 */
public void setListenerIndex(int index)
{
    this.listenerIndex = index;
} //end of setListenerIndex() method

/*
 * Gets the encapsulated event
 * @return Return the event encapsulated by this object
 */
public Object getEvent()
{
    return this.event;
} //end of getEvent() method

/*
 * Gets the class type of encapsulated event.
 * @return Returns class type of encapsulated event.
 */
public Class getEventClass()
{
    return this.eventClass;
} //end of getEventClass() method

/*
 * Gets the talker of encapsulated event.
 * @return Returns talker of the event.
 */
public Object getTalker()
{
    return this.talker;
} //end of getTalker() method

```

```

/*
 * Gets the class type of talker.
 * @return Returns class type of talker.
 */
public Class getTalkerClass()
{
    return this.talkerClass;
} //end of getTalkerClass() method

/*
 * Gets the priority of event
 * @return Returns event priority
 */
public int getEventPriority()
{
    return this.eventPriority;
} //end of getEventPriority() method

/*
 * Gets the priority of talker
 * @return Returns talker priority
 */
public int getTalkerPriority()
{
    return this.talkerPriority;
} //end of getTalkerPriority() method

/*
 * Gets listener's index for this event.
 * @return Returns the listener's index for this event.
 */
public int getListenerIndex()
{
    return this.listenerIndex;
} //end of getListenerIndex() method

/*
 * Gets the instantiation time of this event. Time is system time in milliseconds
 * @return Returns instantiation time of this event.
 */
public long getTimestamp()
{
    return this.timestamp;
} //end of getTimestamp() method

} //end of ChannelEvent

```

```
package org.saamnet.channel;
```

```
/*
```

```
* Event filtering can be used to multicast the events to interested listeners  
* instead of broadcasting to all listeners. Multicasting events reduce total  
* delivery time spent by channel for an event. As a result, channel throughput  
* will increase. On the listener side, channel listeners have to check all coming  
* events to choose the desired ones without filtering. It can be guaranteed that  
* channel listeners get the events that they interest when event filtering is  
* employed. ChannelFilter interface is designed for providing filtering capability  
* to channel. It has one simple method needed to implement.
```

```
*/
```

```
public interface ChannelFilter
```

```
{
```

```
/*
```

```
* Checks that if an event is acceptable or not.  
* @param event Event to be checked.  
* @return Returns true if event is acceptable, else false.
```

```
*/
```

```
public boolean isAccepted (ChannelEvent event);
```

```
}
```

```
package org.saamnet.channel;
```

```
/*
```

```
* All channel listeners must implement ChannelListener interface to  
* register with a channel and get events from this channel. ChannelListener  
* interface contains just one method:  
*
```

```
*     public void receiveEvent (ChannelEvent event )
```

```
* This method is called by channel to deliver an event to registered channel  
* listeners. This method should be synchronized to guarantee that no more than  
* one channel dispatcher thread is allowed when a channel listener registers with  
* multiple channels.
```

```
*/
```

```
public interface ChannelListener
```

```
{
```

```
    /**
```

```
     * This method is called by the Channel to pass event notifications  
     * to the Objects that implement this interface.
```

```
     * @param event Channel event
```

```
     */
```

```
    public void receiveEvent(ChannelEvent event);
```

```
}
```

```

package org.saamnet.channel;

/*
 * Channel scheduler allows channel to buffer and dispatch events to multiple
 * listeners in a specific order. Channel events can be categorized in accordance
 * with their talker, event types and priorities. This categorization is used to
 * specify the event delivery order of channel.Channel scheduler buffers and orders
 * channel events according to its scheduling policy. Channel dispatches events to
 * listeners by pulling from scheduler. Channel scheduler is implemented as an
 * interface to provide functional flexibility for channel. Channel forwards events
 * to scheduler by calling push method. On the other hand, channel dispatches events
 * by extracting from scheduler via pull method. Channel package includes three built-in
 * schedulers, which implement ChannelScheduler interface. Developers can implement
 * and deploy their own schedulers that meet specific requirements.
 */
public interface ChannelScheduler
{
    /*
     * Inserts an event to scheduler
     * @param event Event to be pushed
     * @return Returns true if insertion is succesful, else false due to buffer overflow
     */
    public boolean push (ChannelEvent event);

    /*
     * Extracts an event from scheduler
     * @return Returns a channel event.
     */
    public ChannelEvent pull ();
}
// End of ChannelScheduler

```

```

package org.saamnet.channel;

import java.util.*;

/**
 * This class provides a data structure for channel to keep the channel
 * listeners and their properties. Listener itself, its filter and priority
 * are encapsulated by this class as a channel listener item.
 * This class also provides self event dispatching capability to channel
 * listeners. If the execution of a listener's "receiveEvent" method takes
 * long time, this will cause the other listeners to get events late. To prevent
 * this handicap, channel offers self-dispatching to these listeners which need
 * more excution time. Channel ask this class to create a new dedicated thread
 * for dispatching events which belongs to the encapsulated listener.
 * After starting the self-dispatching, channel will not call listener's recieveEvent method
 * anymore.Channel will put the coming events into the listener's event queue,
 * self-dispatching thread will deliver the events in this queue.
 */

public class ChannelListenerItem implements Runnable
{
    /**
     * Channel listener to be encapsulated by this class
     */
    private ChannelListener listener;

    /**
     * Keeps the priority of this listener
     */
    private int priority ;

    /**
     * Keeps the filters of this listener
     */
    private Vector listenerFilters;

    /**
     * Indicates that this listener has a filter or not
     */
    private boolean filterFlag = false;

    /**
     * Fifo queue for self-dispatching
     */
    private FIFOQueue listenerQueue;

    /**
     * Indicates the self dispatching

```

```

*/
private volatile boolean selfDispatch = false;

/**
 * Self-dispatching thread.
 * This thread is created when self dispatching is activated.
 */
private Thread selfThread;

/**
 * Indicates the status of self dispatching thread; asleep or awake.
 */
private boolean selfThreadAsleep;

/**
 * Channel listener to be encapsulated by this class
 */
private Object theSelfThreadLock;

/**
 * Indicates that self-dispatching thread will continue to run or it
 * will be ended.
 */
private volatile boolean stopSelfDispatch;

/**
 * Indicates that a suspend request was received for the self-dispatching
 * thread.
 */
private boolean suspendRequest = false;

/**
 * Indicates that a resume request was received for the self-dispatching
 * thread
 */
private boolean resumeRequest = false;

/**
 * Indicates that a stop request was received for the self-dispatching
 * thread.
 */
private boolean stopRequest = false;

/**
 * Constructs a channel listener item by encapsulating the channel
 * listener and default properties.
 * @param listener a Channel listener
 */
ChannelListenerItem (ChannelListener listener)
{

```



```

        this.listener = listener;
        this.priority = Channel.DEFAULT_LISTENER_PRIORITY;
    }// end of constructor

/**
 * Constructs a channel listener item by encapsulating the channel
 * listener and its filter object.
 * @param listener a Channel listener
 * @param eventFilter listener's filter
 */
    ChannelListenerItem (ChannelListener listener,ChannelFilter eventfilter)
    {
        this.listener = listener;
        this.priority = Channel.DEFAULT_LISTENER_PRIORITY;
        this.listenerFilters = new Vector();
        listenerFilters.add(eventfilter);
        this.filterFlag = true;
    }// end of constructor

/**
 * Constructs a channel listener item by encapsulating the channel
 * listener and its priority.
 * @param listener a Channel listener
 * @param priority listener's priority
 */
    ChannelListenerItem (ChannelListener listener, int priority)
    {
        this.listener = listener;
        this.priority = priority;
    }// end of constructor

/**
 * Constructs a channel listener item by encapsulating the channel
 * listener and its filter object and priority.
 * @param listener a Channel listener
 * @param eventFilter listener's filter
 * @param priority listener's priority
 */
    ChannelListenerItem (ChannelListener listener,
                        ChannelFilter eventfilter ,
                        int priority)
    {
        this.listener = listener;
        this.priority = priority;
        this.listenerFilters = new Vector();
        listenerFilters.add(eventfilter);
        this.filterFlag = true;
    }// end of constructor

/**
 * Sets this listener's priority

```

```

* @param Listener's priority
*/
void setListenerPriority(int priority)
{
    this.priority = priority;
} //end of setListenerPriority() method

/**
* Sets suspend request flag of this listener
* @param order true if there is a request , false otherwise
*/
void setSuspendRequest(boolean order)
{
    suspendRequest = order;
} //end of setSuspendRequest() method

/**
* Sets resume request flag of this listener
* @param order true if there is a request , false otherwise
*/
void setResumeRequest(boolean order)
{
    resumeRequest = order;
} //end of setResumeRequest() method

/**
* Sets stop request flag of this listener
* @param order true if there is a request , false otherwise
*/
void setStopRequest(boolean order)
{
    stopRequest = order;
} //end of setStopRequest() method

/**
* Tests if there is a suspend request or not.
* @return Returns true if a request was made, false otherwise
*/
boolean hasSuspendRequest()
{
    return suspendRequest;
} //end of hasSuspendRequest() method

/**
* Tests if there is a resume request or not.
* @return Returns true if a request was made, false otherwise
*/
boolean hasResumeRequest()
{
    return resumeRequest;
} //end of hasResumeRequest() method

```

```

/**
 * Tests if there is a stop request or not.
 * @return Returns true if a request was made, false otherwise
 */
boolean hasStopRequest()
{
    return stopRequest;
} //end of hasStopRequest() method

/**
 * Tests if the self-dispatching queue is empty or not.
 * @return Returns true if the self-dispatching queue is empty,
 * false otherwise
 */
boolean isSelfDispatchQueueEmpty()
{
    if(isSelfDispatching())
    {
        return listenerQueue.isEmpty();
    }
    else
    {
        return false;
    }
} //end of isSelfDispatchQueueEmpty() method

/**
 * Sets self dispatch flag of this listener.
 * This method is used by channel to set this flag externally.
 * @param setting true for self-dispatching, false otherwise
 */
void setSelfDispatch(boolean setting)
{
    selfDispatch = setting;
} //end of setSelfDispatch() method

/**
 * Starts self-dispatching for this listener.
 * Self-dispatching thread and a fifo queue is created.
 * After starting the self-dispatching, channel gives up delivering
 * the events by calling listener's "receiveEvent" method and then
 * it starts to put the event into self-dispatching queue for self
 * delivery.
 */
void startSelfDispatching()
{
    if(! isSelfThreadAlive())
    {
        selfThread = new Thread(this);
        theSelfThreadLock = new Object();
    }
}

```

```

    listenerQueue = new FIFOQueue(1000);
    selfThreadAsleep = false;
    selfDispatch = true;
    stopSelfDispatch = false;
    selfThread.start();
}
} //end of startSelfDispatching() method

/**
 * Stops self dispatching for this listener.
 * This will end the self-dispatching thread.
 * Channel waits appropriate situation(when self-dispatch queue is empty)
 * to call this method for a graceful shutdown.
 */
void stopSelfDispatching()
{
    stopSelfDispatch = true;
    if(selfThreadAsleep )
    {
        synchronized(theSelfThreadLock)
        {
            theSelfThreadLock.notify();
        }
    }
} //end of stopSelfDispatching() method

/**
 * Tests if self-dispatching thread is alive or not
 * @return Returns true if it is alive, false otherwise.
 */
public boolean isSelfThreadAlive()
{
    if(selfThread==null) return false;
    return selfThread.isAlive();
} //end of isSelfThreadAlive() method

/**
 * Channel uses this method to put an event into self-dispatching
 * queue. This method notifies the self-dispatching thread if it is
 * asleep.
 * @param event event to be enqueued
 */
void putToListenerQueue(ChannelEvent event)
{
    listenerQueue.enqueue(event);
    if(selfThreadAsleep )
    {
        synchronized(theSelfThreadLock)
        {
            theSelfThreadLock.notify();
        }
    }
}

```

```

    }
} //end of putToListenerQueue() method

/**
 * Tests if the self-dispatching is active or not
 * @return Returns true if self-dispatching is active, false otherwise
 */
boolean isSelfDispatching()
{
    return selfDispatch;
} //end of isSelfDispatching() method

/**
 * Tests if the event is acceptable by this listener or not.
 * If there is no filter object which is associated with this listener,
 * it is always acceptable.
 * Concatenation of filters is also possible. Filters of a listener are
 * stored in a vector. An event is only accepted if it passes all filters.
 * @param event Event to be tested.
 * @return Returns true if the event is acceptable, false otherwise
 */
boolean isAcceptedEvent(ChannelEvent event)
{
    if(filterFlag)
    {
        Iterator i = listenerFilters.iterator();
        while(i.hasNext())
        {
            if( ! ((ChannelFilter) i.next()).isAccepted(event) )
            {
                return false;
            }
        }
        return true;
    }
    else
    {
        return true;
    }
} //end of isAcceptedEvent() method

/**
 * Add a new filter to listener filter vector
 * @param newFilter New filter object to be added
 */
void addFilter(ChannelFilter newFilter)
{
    if(filterFlag)
    {
        listenerFilters.add(newFilter);
    }
}

```

```

    else
    {
        if(listenerFilters == null)
        {
            listenerFilters = new Vector();
        }
        listenerFilters.add(newFilter);
        filterFlag=true;
    }
} //end of addFilter() method

/**
 * Remove an old filter from listener filter vector.
 * @param newFilter Old filter object to be removed
 */
void removeFilter(ChannelFilter oldFilter)
{
    if(filterFlag)
    {
        listenerFilters.remove(oldFilter);
    }
} //end of removeFilter() method

/**
 * Remove all filters that are belong to this listener.
 * After this, all events will be accepted until adding
 * adding a new filter
 */
void removeAllFilters()
{
    if(filterFlag)
    {
        listenerFilters.removeAllElements();
        filterFlag = false;
    }
} //end of removeAllFilters() method

/**
 * Gets the channel listener which is encapsulated by this class
 * @return Returns channel listener
 */
public ChannelListener getListener()
{
    return listener;
} //end of getListener() method

/**
 * Gets this listener's priority
 * @return Listener's priority
 */
public int getListenerPriority()

```

```

{
    return priority;
} //end of getListenerPriority() method

/**
 * Gets the filter which belongs to this listener
 * @return Returns the listener's filter
 */
public Vector getListenerFilters()
{
    return (Vector) listenerFilters.clone();
} //end of getListenerFilters() method

/**
 * Overrides the Object's "toString()" method.
 * @return Returns string representation of this listener.
 */
public String toString()
{
    return listener.toString();
} //end of toString() method

/**
 * The "run" method for self-dispatching thread. If a channel listener ask
 * a new thread due to its "recieveEvent" method which takes long time, this
 * class creates a new thread on behalf of this listener and the other
 * listeners in the channel. The other listeners will not suffer from this
 * listener's long execution after self-dispatching.
 * Channel controls the starting,suspending,resuming and stopping of this
 * self-dispatching thread.
 */
public void run()
{
    if(Thread.currentThread() != selfThread)
    {
        throw new RuntimeException( " This is a self-running object, this"+
            " method can be called by only object itself");
    }

    while(!stopSelfDispatch )
    {
        if(listenerQueue.isEmpty())
        {
            try
            {
                synchronized(theSelfThreadLock)
                {
                    selfThreadAsleep = true;
                    theSelfThreadLock.wait();
                    selfThreadAsleep = false;
                }
            }
        }
    }
}

```

```
    }
    catch(InterruptedException ie)
    {
        ie.printStackTrace();
    }
}
else
{
    ChannelEvent top = (ChannelEvent)listenerQueue.dequeue();
    listener.receiveEvent(top);
}
}
} //end of run() method
} //end of ChannelListenerItem
```



```

package org.saamnet.channel;

import java.util.*;

/**
 * This class is a built-in channel event filter for channel package.
 * It implements Filter interface and provides a high level filtering
 * capability for channel on behalf of the channel listeners.
 * Channel listeners can provide a filter object to the channel when
 * they are registering with the channel. Channel will not deliver the
 * events which is not accepted by this filter.
 * A filter object can be an instance of this ChannelEventFilter class or
 * an instance of another class which implements Filter interface. Developers
 * can build their own low level and detailed filters by using other event
 * properties.
 * This class provides a high level filtering by using event type (the Class
 * that event object belongs), talker type (the Class that the source of event
 * belongs), and talker's name (string which is returned by talker object's
 * toString() method).
 * Channel listener can also stop listening the channel temporarily without
 * unregistration with the channel by using this filter class.
 */

public class ChannelEventFilter implements ChannelFilter
{
    /**
     * Keeps the accepted event types with no talker type or name restriction.
     */
    private Set unrestrictedEventTypes = Collections.synchronizedSet(new HashSet());

    /**
     * Keeps the accepted talker types with no talker type or event type
     * restriction.
     */
    private Set unrestrictedTalkerTypes = Collections.synchronizedSet(new HashSet());

    /**
     * Keeps the accepted talker names
     */
    private Set acceptedTalkerNames = Collections.synchronizedSet(new HashSet());

    /**
     * Keeps the accepted talker types with their accepted names.
     */
    private Vector talkerTypeNamePairs = new Vector ();

    /**
     * Keeps the accepted event types with their accepted talkers' types
     * and names.
     */
    private Vector talkerEventPairs = new Vector ();
}

```

```

/**
 * It is used to stop or restart accepting events
 */
private boolean denyAllFlag = false;

/**
 * This method is used to add an accepted event class without
 * any further restriction.All events which belong to this class
 * are accepted by filter without looking their talker's class or
 * names
 * @param eventType event class to be accepted
 */
public void addAcceptedEvent(Class eventType)
{
    unrestrictedEventTypes.add(eventType);
} // end of addAcceptedEvent() method

/**
 * This method is used to add an accepted talker according to
 * its name. Talker name is obtained from toString() method of object.
 * @param talkerName name of the talker to be accepted
 */
public void addAcceptedTalker(String talkerName)
{
    acceptedTalkerNames.add(talkerName);
} // end of addAcceptedTalker() method

/**
 * This method is used to add an accepted talker class without
 * any further restriction.All events which are coming from this
 * talker class are accepted by filter without looking their talker's
 * names or event classes.
 * @param talkerType talker class to be accepted
 */
public void addAcceptedTalker(Class talkerType)
{
    unrestrictedTalkerTypes.add(talkerType);
} // end of addAcceptedTalker() method

/**
 * This method is used to add an accepted talker class with its
 * name.All events which are coming from this talker class and
 * specified name are accepted by filter without looking their
 * event classes.
 * @param talkerType talker class to be accepted
 * @param talkerName talker's name belongs to accepted talker class
 */
public void addAcceptedTalker(Class talkerType, String talkerName)
{

```

```

    TalkerTypeNamePair newItem = new TalkerTypeNamePair(talkerType,talkerName);
    if(!talkerTypeNamePairs.contains(newItem))
    {
        talkerTypeNamePairs.add(newItem);
    }
} // end of addAcceptedTalker() method

/**
 * This method is used to add an accepted event class with its
 * specified talker class.All events which belongs to this event
 * class and coming from specified talker class are accepted by filter.
 * @param talkerType talker's class of the accepted pair.
 * @param eventType event's class of the accepted pair.
 */
public void addAcceptedTalkerEventPair( Class talkerType, Class eventType)
{
    TalkerEventTypePair newPair = new TalkerEventTypePair(talkerType,eventType);
    if( ! talkerEventPairs.contains(newPair) )
    {
        talkerEventPairs.add(newPair);
    }
} // end of addAcceptedTalkerEventPair() method

/**
 * This method is used to add an accepted event class with its
 * specified talker class and name.All events which belongs to this event
 * class and coming from specified talker class and name are accepted
 * by filter.
 * @param talkerType talker's class of the accepted pair.
 * @param talkerName talker's name of the accepted pair
 * @param eventType event's class of the accepted pair.
 */
public void addAcceptedTalkerEventPair( Class talkerType,
                                         String talkerName,
                                         Class eventType)
{
    TalkerEventTypePair newPair = new TalkerEventTypePair(talkerType,
                                                         talkerName,
                                                         eventType);
    if( ! talkerEventPairs.contains(newPair) ){
        talkerEventPairs.add(newPair);
    }
} // end of addAcceptedTalkerEventPair() method

/**
 * This method is used to remove the accepted event classes which were added
 * by "addAcceptedEventType" method.
 * @param eventType event class to be removed
 */
public void removeAcceptedEvent(Class eventType)

```

```

{
    if(unrestrictedEventTypes.contains (eventType))
    {
        unrestrictedEventTypes.remove(eventType);
    }
    else
    {
        System.out.println("There is no such event-type");
    }
}
} // end of removeAcceptedEvent() method

```

```

/**
 * This method is used to remove the accepted talker names which were added
 * by "addAcceptedTalkerName" method.
 * @param talker name to be removed
 */
public void removeAcceptedTalker(String talkerName)
{
    if(acceptedTalkerNames.contains (talkerName))
    {
        acceptedTalkerNames.remove(talkerName);
    }
    else
    {
        System.out.println("There is no such talker name");
    }
}
} // end of removeAcceptedTalker() method

```

```

/**
 * This method is used to remove the accepted talker classes which were added
 * without a specified talker's name.
 * @param talkerType talker class to be removed
 */
public void removeAcceptedTalker(Class talkerType)
{
    if(unrestrictedTalkerTypes.contains (talkerType))
    {
        unrestrictedTalkerTypes.remove(talkerType);
    }
    else{
        System.out.println("There is no such talker-type");
    }
}
} // end of removeAcceptedTalker() method

```

```

/**
 * This method is used to remove the accepted talker classes which were added
 * with a specified talker's name.
 * @param talkerType talker class to be removed
 * @param talkerName talker's name

```

```

*/
public void removeAcceptedTalker(Class talkerType, String talkerName)
{
    TalkerTypeNamePair removedItem = new TalkerTypeNamePair ( talkerType,
                                                                talkerName);
    if(talkerTypeNamePairs.contains (removedItem))
    {
        unrestrictedTalkerTypes.remove(removedItem);
    }
    else
    {
        System.out.println("There is no such talker-type");
    }
}
} // end of removeAcceptedTalker() method

/**
 * This method is used to remove an accepted event class with its
 * specified talker class.
 * @param talkerType talker class of the pair to be removed
 * @param eventType event class of the pair to be removed
 */
public void removeAcceptedTalkerEventPair( Class talkerType, Class eventType)
{
    TalkerEventTypePair removedPair = new
TalkerEventTypePair(talkerType,eventType);
    if( talkerEventPairs.contains(removedPair) )
    {
        talkerEventPairs.remove(removedPair);
    }
    else
    {
        System.out.println("There is no such event-talker pair");
    }
}
} // end of removeAcceptedTalkerEventPair() method

/**
 * This method is used to remove an accepted event class with its
 * specified talker class and name.
 * @param talkerType talker class of the pair to be removed
 * @param talkerName talker name of the pair to be removed
 * @param eventType event class of the pair to be removed
 */
public void removeAcceptedTalkerEventPair( Class talkerType,
                                             String talkerName,
                                             Class eventType)
{
    TalkerEventTypePair removedPair = new TalkerEventTypePair(talkerType,
                                                                talkerName,
                                                                eventType);
    if( talkerEventPairs.contains(removedPair) )
    {

```

```

        talkerEventPairs.remove(removedPair);
    }
    else
    {
        System.out.println("There is no such event-talker pair");
    }
} // end of removeAcceptedTalkerEventPair() method

/**
 * Filter denies all coming events, regardless of these events are acceptable
 * or not. This method provides listeners to stop accepting events from channel
 * temporarily without unregistration with channel.
 */
public void suspendAccepting()
{
    denyAllFlag = true ;
} // end of suspendAccepting() method

/**
 * Returns the filter to its normal operation after "suspendAccepting"
 * method. Filter begins to accept events according to its existing
 * settings.
 */
public void resumeAccepting()
{
    denyAllFlag = false ;
} // end of resumeAccepting() method

/**
 * Tests if this event's properties is in the filter's database or not
 * @param event event to be tested
 * @return Returns true if filter's database contains this event type
 * or talker type, falsed otherwise.
 */
private boolean contains(ChannelEvent event)
{
    boolean isContain = false;
    Class talkerType = event.getTalkerClass();
    Class eventType = event.getEventClass();
    if(unrestrictedTalkerTypes.contains(talkerType)||
        unrestrictedEventTypes.contains(eventType) )
    {
        isContain = true;
    }
    else
    {
        if(talkerTypeNamePairs.contains( new TalkerTypeNamePair(talkerType,
            (event.getTalker()).toString()) )
        {
            isContain = true;
        }
    }
}

```

```

    }
    else
    {
        if(talkerEventPairs.contains( new TalkerEventTypePair(talkerType,eventType))
            ||talkerEventPairs.contains( new TalkerEventTypePair(talkerType,
                (event.getTalker()).toString(),
                eventType)))
        {
            isContain = true;
        }
        else{
            isContain = false;
        }
    }
}
return isContain;
} // end of contains() method

/**
 * Remove all accepted types and clears filter's database
 */
public void removeAll()
{
    unrestrictedEventTypes.clear();
    unrestrictedTalkerTypes.clear();
    talkerTypeNamePairs.clear();
    talkerEventPairs.clear();
} // end of removeAll() method

/**
 * Test if this event is acceptable by filter or not. If this event is
 * not acceptable, channel will not deliver the event to the listener
 * which is the owner of this filter.
 * @param event event to tested
 * @return Returns true if this event is acceptable, false otherwise
 */
public boolean isAccepted(ChannelEvent event)
{
    if(! denyAllFlag && this.contains(event) )
    {
        return true;
    }
    else
    {
        return false;
    }
} // end of isAccepted() method

/**
 * Data structure for creating talker-event type pairs. This structure
 * defines a event type with specific talker's type and name.

```

```

*
*/
private class TalkerEventTypePair{

    /** Keeps talker type */
    private Class talkerType;

    /** Keeps event type */
    private Class eventType;

    /** Keeps talker's name */
    private String talkerName;

    /** Default talker's name if it is not specified */
    private final static String defaultName = "ANY";

    /** Constructs a pair with default talker's name
     * @param talkerType talker type
     * @param eventType event type
     */
    TalkerEventTypePair(Class talkerType, Class eventType)
    {
        this.talkerType = talkerType;
        this.eventType = eventType;
        this.talkerName = defaultName;
    } //end of constructor

    /** Constructs a pair with specified event type, talker type and
     * talker's name
     * @param talkerType talker type
     * @param talkerName talker's name
     * @param eventType event type
     */
    TalkerEventTypePair(Class talkerType, String talkerName, Class eventType)
    {
        this.talkerType = talkerType;
        this.eventType = eventType;
        this.talkerName = talkerName;
    } //end of constructor

    /** Overrides Object's equal method for equality checking
     * @param obj Object to be compared
     * @return Returns true if it is equal, false otherwise
     */
    public boolean equals(Object obj)
    {
        if(talkerType.equals( ((TalkerEventTypePair)obj).talkerType ) &&
            eventType.equals( ((TalkerEventTypePair)obj).eventType) &&
            talkerName.equals( ((TalkerEventTypePair)obj).talkerName))
        {
            return true;
        }
    }
}

```



```

    }
    else
    {
        return false;
    }
} //end of () method
} //end of TalkerEventTypePair

/**
 * Data structure for creating a talker type with its name.
 */
private class TalkerTypeNamePair{

    /** Keeps talker type */
    private Class talkerType;

    /** keeps talker's name */
    private String talkerName;

    /** Constructs a talker type with its name
     * @param talkerType talker type
     * @param name talker name
     */
    TalkerTypeNamePair(Class talkerType, String name)
    {
        this.talkerType = talkerType;
        this.talkerName = name;
    } //end of constructor

    /** Overrides Object's equal method for equality checking
     * @param obj Object to be compared
     * @return Returns true if it is equal, false otherwise
     */
    public boolean equals(Object obj)
    {
        if(talkerType.equals( ((TalkerEventTypePair)obj).talkerType ) &&
            talkerName.equals( ((TalkerEventTypePair)obj).talkerName))
        {
            return true;
        }
        else
        {
            return false;
        }
    } // end of equals() method
} //end of TalkerTypeNamePair
} //end of ChannelEventFilter

```

```

package org.saamnet.channel;

/*
 * FIFOScheduler class simply buffers arriving events into a First In First Out
 * (FIFO) queue. Events are served by their arriving orders. FIFOScheduler is
 * default scheduler for channel when another scheduler is not specified.
 */

public class FIFOScheduler implements ChannelScheduler
{

    /*
     * Stores channel events
     */
    private FIFOQueue queue;

    /**
     * default size of the queue
     */
    public static final int DEFAULT_SIZE = 1000;

    /*
     * Keeps the max queue size
     */
    private int maxSize = DEFAULT_SIZE ;

    /*
     * Constructs an FIFOScheduler with default queue size
     */
    public FIFOScheduler ()
    {
        this.queue = new FIFOQueue (maxSize);
    }// end of constructor

    /*
     * Constructs an FIFOScheduler with specified queue size
     */
    public FIFOScheduler (int queueSize)
    {
        maxSize = queueSize;
        this.queue = new FIFOQueue (maxSize);
    }// end of constructor

    /*
     * Enqueue a channel event into FIFO queue
     * @param event Channel event to be enqueued
     */
    public boolean push(ChannelEvent event)
    {
        if (getSize() >= maxSize)
        {

```

```

        System.out.println("FIFOScheduler dropped a packet due "+
            "to buffer overflow");
        return false;
    }
    else
    {
        queue.enqueue(event);
        return true;
    }
} //end of push() method

/*
 * Dequeue a channel event from FIFO queue
 * @return Returns the channel event at the top of queue
 */
public ChannelEvent pull ()
{
    return (ChannelEvent) queue.dequeue();
} //end of pull() method

/*
 * Gets the number of events in the queue
 * @return Returns the number of channel events in the queue.
 */
public int getSize()
{
    return queue.getSize();
} //end of getSize() method

} //end of FIFOScheduler

/**
 * this class implements a FIFO queue
 * Jan2000 [Uysal, Xie] - Created
 */
class FIFOQueue
{
    /**
     * default size of the queue
     */
    public static final int DEFAULT_SIZE = 1000;

    /**
     * max size of the queue
     */
    private int maxSize;

    /**
     * Keeps lost items

```

```

*/
private int lossCount;

/**
 * current size of the queue
 */
private int size;

private FifoQueueItem first, last, current;

/**
 * parameterless constructor of the class
 */
public FIFOQueue()
{
    maxSize = this.DEFAULT_SIZE;
    size = lossCount = 0;
    first = last = current = null;
} //end constructor

/**
 * parameter constructor
 * @param queuesize int size of the queue
 */
public FIFOQueue(int queuesize)
{
    maxSize = queuesize;
    size = lossCount = 0;
    first = last = current = null;
} //end constructor

/**
 * queues a new packet
 * @param packet Object
 */
public synchronized void enqueue(Object packet)
{
    //If size exceeds the maxSize
    //we need to drop the packet since there is no place in the queue
    if (size >= maxSize)
    {
        lossCount++;
        return;
    }

    size++;

    FifoQueueItem newItem = new FifoQueueItem(packet);

```

```

    if (first == null)
    {
        first = last = newItem;
    }
    else
    {
        current = last;
        current.next = newItem;
        newItem.previous = current;
        last = newItem;
    } //end else
} //end enqueue()

/**
 * removes the first packet from the queue
 * @return Object
 */
public synchronized Object dequeue()
{
    if (first != null)
    {
        size--;
        current = first;
        if (first.next == null)
        { //item is also the last
            first = last = null;
        }
        else
        {
            first = current.next;
            first.previous = null;
        }
    }

    return current.data;
} //end if
return null;
} //end method dequeue()

/**
 * Returns the Object at head of queue.
 * Under heavy loading, <code>peek</code> may find the queue empty
 * and the return will cause a NullPointerException to be thrown.
 * @return Object
 */
public Object peek()
{
    try

```

```

    {
        return first.data;
    }
    catch(NullPointerException npe)
    {
        System.out.println("****[DIAG MSG]*** FIFOQueue.peek(): No enqueued packet
found. ****");
        throw new NullPointerException();
    }
} //end method peek()

```

```

/**
 * returns the size of the queue
 * @return int
 */
public int getSize()
{
    return size;
} //end method getSize()

```

```

/**
 * returns true if the queue is empty
 * @return boolean
 */
public synchronized boolean isEmpty()
{
    return (size == 0);
} //end isEmpty()

```

```

/**
 * returns the max size of the queue
 * @return int
 */
public int getMaxQueueSize()
{
    return this.maxSize;
} //end method getMaxQueueSize()

```

```

/**
 * changes the max queue size setting
 * @param newSize int size to be set
 */
public synchronized void setMaxQueueSize(int newSize)
{
    //max size should be at least 1
    if (newSize < 1)
    {
        return;
    }
}

```

```

    this.maxSize = newSize;
} //end method setQueueSize()

/**
 * returns the int packet loss count of the queue
 * @return int
 */
public int getLossCount()
{
    return this.lossCount;
} //end method getLossCount()

/**
 * resets the packet loss count to 0
 */
public synchronized void clearLossCount()
{
    this.lossCount = 0;
} //end method clearLossCount()

/**
 * return a string representation of the queue
 * @return String
 */
public String toString()
{
    return "FIFO Queue";
} //end method toString()

} //end class FIFOQueue

class FifoQueueItem
{
    public Object data;
    public FifoQueueItem next, previous;

    public FifoQueueItem(Object packet)
    {
        data = packet;
        next = null;
        previous = null;
    } //end constructor
} //end class QueueItem

//end file FIFOQueue

```

```

package org.saamnet.channel;

/*
 * It offers a fair serving distribution between channel talkers.
 * A priority queue is created for each talker. Arriving events are buffered
 * in appropriate queue in accordance with their talkers. Scheduler extracts
 * events from talker queues based on a round-robin algorithm.
 */
import java.util.*;

public class PerTalker_RR_Scheduler implements ChannelScheduler {

    /*
     * Keeps channel talkers
     */
    private Vector talkerQueues;

    /*
     * Keeps the channel which this scheduler is used.
     */
    private Channel parentChannel;

    /**
     * default size for talker queues
     */
    public static final int DEFAULT_SIZE = 1000;

    /*
     * Keeps the max queue size for talker queues
     */
    private int maxSize = DEFAULT_SIZE ;

    /*
     * Indicates the current talker to be served
     */
    private int currentTalkerIdx = 0;

    /*
     * Constructs a PerTalker_RR_Scheduler with default talker queue size
     */
    public PerTalker_RR_Scheduler()
    {
        this.talkerQueues = new Vector();
    } // end of constructor

    /*
     * Constructs a PerTalker_RR_Scheduler with specified queue size
     * @param q_size The size of talker queues
     */
    public PerTalker_RR_Scheduler(int q_Size)

```



```

{
    this.talkerQueues = new Vector();
    this.maxSize = q_Size;
} // end of constructor

/*
 * Sets the parent channel that this scheduler belongs.
 * This scheduler will be active after setting the parent channel.
 * Channel automatically calls this method if its scheduler type
 * is PerTalker_RR_Scheduler
 * @param parent Parent channel of this scheduler
 */
public void setParentChannel(Channel parent)
{
    this.parentChannel = parent;
    updateTalkerQueues();
} //end of setParentChannel() method

/*
 * This methods updates the talkers vector of scheduler. Updating necessary if
 * new talker is added to channel.
 */
private void updateTalkerQueues()
{
    Vector copy = (Vector) ( parentChannel.getTalkers() ).clone();
    Enumeration e = copy.elements();
    while( e.hasMoreElements() )
    {
        Object tempElement = e.nextElement();
        TalkerQueue temp = new TalkerQueue (tempElement,maxSize );
        if( ! talkerQueues.contains(temp) )
        {
            talkerQueues.add (temp);
        }
    }
} //end of updateTalkerQueues() method

/*
 * Inserts an event to appropriate talker queue by checking
 * event's talker.
 * @param event Channel event
 * @return Returns true if insertion is successful,returns false
 * if a talker queue, which is matching with talker of this event,
 * could not be found.
 */
private boolean insert(ChannelEvent event)
{
    boolean result = false ;
    TalkerQueue temp = new TalkerQueue (event.getTalker(),maxSize );
    if( ! talkerQueues.contains(temp) )

```

```

    {
        updateTalkerQueues();
    }
    Enumeration e = talkerQueues.elements();
    while( e.hasMoreElements() )
    {
        TalkerQueue tempElement = (TalkerQueue) e.nextElement();
        if( (tempElement.getTalker()).equals(event.getTalker()) )
        {
            if(tempElement.enqueue(event))
            {
                result=true;
            }
            else
            {
                result=false;
            }
            break;
        }
    }
    return result;
} //end of insert() method

/*
 * Pushes an event to appropriate talker queue by using "insert"
 * method. If insertion is failed, updates talker queues and tries again.
 * Channel quarantees that only registered talkers can be talk on channel.
 * @param event Channel event
 */
public boolean push(ChannelEvent event)
{
    if(parentChannel == null) return false;
    return insert(event);
} //end of push() method

/*
 * Extracts events from talker queues based on a round-robin algorithm.
 * @return Returns a channel event.
 */
public ChannelEvent pull()
{
    int index = currentTalkerIdx;
    if( talkerQueues.isEmpty() ) return null;
    while( ((TalkerQueue) talkerQueues.elementAt(index) ).isEmpty() )
    {
        index++;
        if( index == talkerQueues.size() )
        {
            index=0;
        }
    }
    if( index == currentTalkerIdx )

```

```

    {
        return null;
    }
}

ChannelEvent event = ( (TalkerQueue) talkerQueues.elementAt(index) ).dequeue();
index ++;
if( index == talkerQueues.size() )
{
    currentTalkerIdx=0;
}
else {
    currentTalkerIdx=index;
}
return event;
} //end of pull() method

/*
 * Gets the total number of events in talker queues
 * @return Returns the number of channel events in the queue.
 */
public int getTotalSize()
{
    int totalSize = 0;
    Enumeration e = talkerQueues.elements();
    while( e.hasMoreElements() )
    {
        TalkerQueue tempElement = (TalkerQueue) e.nextElement();
        totalSize +=tempElement.getSize();
    }
    return totalSize;
} //end of getTotalSize() method

} //end of PerTalker_RR_Scheduler

/*
 * This class is a data structure to keep channel talker and its event queue.
 */
class TalkerQueue
{
    /*
     * Keeps talker object.
     */
    private Object talker;

    /*
     * Keeps the max queue size for talker queue.
     */
    private int max;

```

```

/*
 * Stores events, which are belong to the talker.
 * In here Channel Priority Scheduler is used as a priority.
 * PriorityScheduler can be also used an stable priority queue.
 */
private PriorityScheduler talkerPriorityQueue;

/*
 * Constructs an talker queue associated with a channel talker
 * @param talker Channel talker
 * @param queueSize maximum size of talker's event queue
 */
TalkerQueue (Object talker, int queueSize)
{
    this.talker=talker;
    max = queueSize;
    this.talkerPriorityQueue= new PriorityScheduler(max);
} //end of constructor

/*
 * Enqueues an event to talker's event queue
 * @param event Channel event
 */
public boolean enqueue(ChannelEvent event)
{
    //If size exceeds the maxSize
    //we need to drop the packet since there is no place in the queue
    if (talkerPriorityQueue.getSize() >= max)
    {
        System.out.println("PerTalkerRoundRobinScheduler dropped a packet due "+
            "to talker queue overflow "+
            "\nTalker Name : "+ talker);
        return false;
    }
    else
    {
        talkerPriorityQueue.push(event);
        return true;
    }
} //end of enqueue() method

/*
 * Dequeues an event from talker's event queue
 * @return Returns a channel event
 */
public ChannelEvent dequeue()
{
    return (ChannelEvent) talkerPriorityQueue.pull();
} //end of dequeue() method

/*

```

```

    * Enqueues an event to talker's event queue
    * @return Returns the talker object encapsulated by this class
    */
    public Object getTalker()
    {
        return this.talker;
    } //end of getTalker() method

    /*
    * Tests if the talker's event queue is empty or not
    * @return Returns true if talker queue is empty, else false
    */
    public boolean isEmpty()
    {
        return talkerPriorityQueue.isEmpty();
    } //end of isEmpty() method

    /*
    * Gets the number of events in the talker's event queue
    * @return Returns the number of channel events in the talker's event queue.
    */
    public int getSize()
    {
        return talkerPriorityQueue.getSize();
    } //end of getSize() method

    /*
    * Overrides Object's "equals" method
    */
    public boolean equals(Object obj)
    {
        if( this.talker == obj)
        {
            return true;
        }
        else{
            return false;
        }
    } //end of equals() method

} //end of TalkerQueue

```

```

package org.saamnet.channel;

import com.objectspace.jgl.PriorityQueue;
import com.objectspace.jgl.BinaryPredicate;

/*
 * Events are buffered in a priority queue in accordance with their priority.
 * Events with highest priority will serve first. This scheduler can be used
 * to give channel delivery precedence to particular event types.
 */

public class PriorityScheduler implements ChannelScheduler
{

    /*
     * Stores channel events
     */
    private PriorityQueue queue;

    /*
     * PriorityQueue is not stable(uses heap sort).This index
     * is used to make PriorityQueue stable
     */
    private int heapIndex = Integer.MIN_VALUE;

    /**
     * default size of the queue
     */
    public static final int DEFAULT_SIZE = 1000;

    /*
     *Keeps the max queue size
     */
    private int maxSize = DEFAULT_SIZE ;

    /*
     * Constructs a priority scheduler
     */
    public PriorityScheduler ()
    {
        queue = new PriorityQueue ( new ChannelPriorityComparator() );
    }// end of constructor

    /*
     * Constructs a priority scheduler
     */
    public PriorityScheduler (int qSize)
    {
        queue = new PriorityQueue ( new ChannelPriorityComparator() );
        maxSize = qSize;
    }
}

```

```

} // end of constructor

/*
 * Insert a channel event to priority queue
 * @param event Channel event
 */
public boolean push(ChannelEvent event)
{
    //If size exceeds the maxSize
    //we need to drop the packet since there is no place in the queue
    if (getSize() >= maxSize)
    {
        System.out.println("Priority Scheduler dropped a packet due "+
            "to queue overflow");
        return false;
    }

    synchronized(queue)
    {
        //reset index if it reaches the max number
        if(heapIndex == Integer.MAX_VALUE)
        {
            heapIndex = Integer.MIN_VALUE;
            PriorityQueue temp = new PriorityQueue(queue);
            queue.clear();
            while(!temp.isEmpty())
            {
                PriorityHeapItem tempItem = (PriorityHeapItem) temp.pop();
                tempItem.setIndex(heapIndex);
                queue.push(tempItem);
                heapIndex++;
            }
        }
        PriorityHeapItem item = new PriorityHeapItem(event,heapIndex);
        heapIndex++;
        queue.push(item);
    }
    return true;
} //end of push() method

/*
 * Extracts an event from priority queue (highest priority)
 * @return Returns the channel event, which has highest priority.
 */
synchronized public ChannelEvent pull ()
{
    if( queue.isEmpty() ) return null;
    return ((PriorityHeapItem) queue.pop()).getEvent();
} //end of pull() method

/*

```

```

* Gets the number of events in the queue
* @return Returns the number of channel events in the queue.
*/
public int getSize()
{
    return queue.size();
} //end of getSize() method

/**
* Return true if I contain no objects.
*/
public boolean isEmpty()
{
    return queue.isEmpty();
} //end of isEmpty() method

/*
* This class is the data structure to encapsulate event,event priority,
* talker priority and order index.
*/
class PriorityHeapItem
{
    /**
    *Keeps the channel event
    */
    private ChannelEvent event;

    /**
    *Keeps the index that is associated with this event
    */
    private int index;

    /**
    *Constructs a priority heap item,which is associated with an index
    *@param channelEvent A channel event
    *@param idx The index to be associated with the event
    */
    PriorityHeapItem(ChannelEvent channelEvent, int idx){
        this.event = channelEvent;
        this.index = idx;
    } //end of constructor

    /**
    *Sets the index for this event
    *@param idx index
    */
    void setIndex ( int idx){
        this.index = idx;
    } //end of setIndex() method

    /**

```



```

    *Gets the event.
    *@return Returns the event that is encapsulated by this heap item
    */
    ChannelEvent getEvent(){
        return event;
    }//end of getEvent() method

    /**
    *Gets the priority of this encapsulated event.
    *@return Returns the priority of this event
    */
    int getEventPriority(){
        return event.getEventPriority();
    }//end of getEventPriority() method

    /**
    *Gets the talker of this encapsulated event.
    *@return Returns the talker of this event
    */
    int getTalkerPriority(){
        return event.getTalkerPriority();
    }//end of getTalkerPriority() method

    /**
    *Gets the index of this item.
    *@return Returns the index of this item.
    */
    int getIndex(){
        return index;
    }//end of getIndex() method

} //end of PriorityHeapItem

/*
 * This class implements BinaryPredicate and is passed to priority
 * queue for comparing queue items.
 */
class ChannelPriorityComparator implements BinaryPredicate
{
    /**
    * Compares two priority queue item
    * @param first First item to be compared
    * @param second Second item to be compared
    * @return Returns true if the first item is less than the second item,
    * else false.
    */
    public boolean execute( Object first, Object second )
    {
        boolean flag = false;
        PriorityHeapItem firstEvent = (PriorityHeapItem)first;

```

```

PriorityHeapItem secondEvent = (PriorityHeapItem)second;
//check talker priority
if(firstEvent.getTalkerPriority() > secondEvent.getTalkerPriority() )
{
    flag = false;
}
else
{
    if(firstEvent.getTalkerPriority() < secondEvent.getTalkerPriority() )
    {
        flag = true;
    }
    else {
        //check event priority if the talker priorities are equal
        if(firstEvent.getEventPriority() > secondEvent.getEventPriority())
        {
            flag = false;
        }
        else
        {
            if(firstEvent.getEventPriority() < secondEvent.getEventPriority())
            {
                flag = true;
            }
            else
            {
                //check index numbers if the event priorities are equal
                if(firstEvent.getEventPriority() == secondEvent.getEventPriority())
                {
                    if(firstEvent.getIndex() < secondEvent.getIndex())
                    {
                        flag = false;
                    }
                    else
                    {
                        flag = true;
                    }
                }
            }
        }
    }
}
return flag;
} //end of () method
} //end of ChannelPriorityComparator
} //end of PriorityScheduler

```

```

package org.saamnet.channel;

/**
 * This interface provides a framework to control and organize all channels
 * and their participants of an application in a centralized manner. It defines
 * five methods to be implemented by such a centralized channel controller.
 */

public interface ChannelManagerAuthority
{

    /**
     * This method allows a channel controller component to specify permissions
     * for talking on all the channels.
     * @param talkerClassName the class name of talker.
     * @param channel_id channel identification number
     * @return Returns true if this talker is authorized on specified channel
     * else false.
     */
    public boolean isTalkerAuthorized (String talkerClassName, int channel_id);

    /**
     * This method allows a channel controller component to specify permissions
     * for listening on all the channels.
     * @param listenerClassName the class name of talker.
     * @param channel_id channel identification number
     * @return Returns true if this listener is authorized on specified channel
     * else false.
     */
    public boolean isListenerAuthorized (String listenerClassName, int channel_id);

    /**
     * This method allows a channel controller component to stipulate the talker
     * priority.
     * @param talkerClassName the class name of talker.
     * @return Return the priority of specified talker
     */
    public int getTalkerPriority(String talkerClassName);

    /**
     * This method allows a channel controller component to stipulate the listener
     * priority.
     * @param listenerClassName the class name of listener.
     * @return Return the priority of specified listener.
     */
    public int getListenerPriority(String listenerClassName);

    /**
     * This method allows a channel controller component to specify proper channel

```

```
* scheduler, which is installed for the channel.  
* @param channel_id channel identification number  
* @return Return proper channel scheduler  
*/  
public ChannelScheduler getSchedulerForChannel(int channel_id);  
  
} //end of ChannelManagerAuthority
```

```

package org.saamnet.channel;

import java.util.Hashtable;
import java.util.Enumeration;
import java.awt.*;
import javax.swing.*;

/**
 * The ChannelManager class was developed to provide a common interface for
 * channel participants in large applications. Basically, this class manages
 * channels by keeping a table of existing channels, creating new channels
 * when it is required, and overriding the Channel class methods with an
 * extra argument to allow the desired channel to be identified simply by
 * its channel_id. A ChannelManager can also be used to control access to
 * channels and enforce predefined priorities for channel talkers and listeners
 * when it is constructed with a channel access authority.
 */
public class ChannelManager
{
    /**
     * Keeps the member channels
     */
    private Hashtable memberChannels;

    /**
     * Keeps the manager authority if it was provided
     */
    private ChannelManagerAuthority authority ;

    /**
     * Indicates that this manager has an authority or not.
     */
    private boolean hasAuthority = false;

    /**
     * Keeps debug statements
     */
    private JTextArea debugText = new JTextArea();

    /**
     * Constructs a ChannelManager object without an authority
     */
    public ChannelManager()
    {
        memberChannels = new Hashtable();
    } // end of constructor

    /**
     * Constructs a ChannelManager object with an authority

```

```

* @param managerAuthority The authority for this ChannelManager object
*/
public ChannelManager(ChannelManagerAuthority managerAuthority)
{
    memberChannels = new Hashtable();
    authority = managerAuthority ;
    hasAuthority = true;
} // end of constructor

/**
 * Talker use this method to push an event to specified channel in form of
 * ChannelEvent object.
 * @param event An instance of ChannelEvent class
 * @param channel_id Channel identification number
 */
public void talk(ChannelEvent event, int channel_id)
{
    Integer key_id = new Integer(channel_id);
    //if it is a unregistered talker, it will be refused by channel.
    if(memberChannels.containsKey(key_id))
    {
        ((Channel) memberChannels.get(key_id)).talk(event);
    }
    else
    {
        System.out.println("There is no channel to talk with this specified id."+
            channel_id);
    }
} //end of talk() method

/**
 * Talkers use this method to push an event to specified channel with its priority.
 * @param talker Event owner
 * @param event Event object
 * @param priority Event priority
 * @param channel_id Channel identification number
 */
public void talk(Object talker, Object event, int priority, int channel_id)
{
    ChannelEvent newEvent = new ChannelEvent(talker,event,priority);
    this.talk(newEvent,channel_id);
} //end of talk() method

/**
 * Talkers use this method to push an event to specified channel without
 * an event priority.Channel encapsulates this event as an ChannelEvent
 * object with default priority (0).
 * @param talker Event owner
 * @param event Event object
 * @param channel_id Channel identification number
 */

```

```

public void talk(Object talker, Object event, int channel_id)
{
    ChannelEvent newEvent = new ChannelEvent(talker,event);
    this.talk(newEvent,channel_id);
} //end of talk() method

/**
 * Adds a new talker to specified channel with priority
 * @param talker a channel talker
 * @param priority talker priority
 * @param channel_id Channel identification number
 */
public void addTalker(Object talker, int priority, int channel_id)
{
    if( this.hasAuthority)
    {
        if(! authority.isTalkerAuthorized(talker.getClass().getName(),channel_id) )
        {
            String warning = "*****" +
                "\n!!!!!!Unauthorized Talker :"+ talker.toString()+
                "\nChannel Id : "+channel_id+
                "\n*****";

            System.out.println(warning);
            debugText.append("\n"+warning);
            return;
        }
        else
        {
            priority = authority.getTalkerPriority(talker.getClass().getName());
        }
    }
    //Only authorized talkers can reach to this point if there is an authority
    Integer key_id = new Integer(channel_id);
    String info = "A new talker was added to specified channel"+
        "\nTalker : "+talker+
        "\nPriority : "+priority+
        "\nChannel : "+ channel_id;
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel) memberChannels.get(key_id)).addTalker(talker,priority);
        debugText.append("\n"+info);
    }
    else
    {
        ChannelScheduler sch;
        if( this.hasAuthority)
        {
            sch = authority.getSchedulerForChannel(channel_id);
        }
        else{
            sch = new FIFOScheduler();
        }
    }
}

```

```

    }
    Channel newChannel = new Channel(channel_id,sch);
    memberChannels.put(key_id,newChannel);
    newChannel.addTalker(talker,priority);
    debugText.append("\n"+info);
    }
} //end of addTalker() method

/**
 * Adds a new talker to specified channel with default priority
 * @param talker A channel talker
 * @param channel_id Channel identification number
 */
public void addTalker(Object talker,int channel_id)
{
    int priority = Channel.DEFAULT_TALKER_PRIORITY;
    this.addTalker(talker,priority,channel_id);
} //end of addTalker() method

/**
 * Adds a new listener to specified channel in the form of ChannelListenerItem.
 * ChannelListenerItem class encapsulates a channel listener with specified
 * priority and a given filter object.
 * @param listenerItem An instance of ChannelListenerItem
 * @param channel_id Channel identification number
 */
public void addListener(ChannelListenerItem listenerItem,int channel_id)
{
    if( this.hasAuthority)
    {
        if(! authority.isListenerAuthorized((listenerItem.getListener()).getClass().getName(),
            channel_id )
        {
            String warning = "*****" +
                "\n!!!!!!Unauthorized Listener :"+(listenerItem.getListener()).toString()+
                "\nChannel Id : "+channel_id+
                "\n*****";
            System.out.println(warning);
            debugText.append("\n"+warning);
            return;
        }
        else
        {
            int prio =
authority.getListenerPriority((listenerItem.getListener()).getClass().getName());
            listenerItem.setListenerPriority(prio);
        }
    }
}
//Only authorized listeners can reach to this point if there is an authority
Integer key_id = new Integer(channel_id);

```



```

String info = "A new listener was added to specified channel"+
    "\nListener : "+listenerItem+
    "\nPriority : "+listenerItem.getListenerPriority()+
    "\nChannel : "+ channel_id;
if(memberChannels.containsKey(key_id) )
{
    ((Channel)memberChannels.get(key_id)).addListener(listenerItem);
    debugText.append("\n"+info);
}
else
{
    ChannelScheduler sch;
    if( this.hasAuthority)
    {
        sch = authority.getSchedulerForChannel(channel_id);
    }
    else
    {
        sch = new FIFOScheduler();
    }

    Channel newChannel = new Channel(channel_id,sch);
    memberChannels.put(new Integer(channel_id),newChannel);
    newChannel.addListener(listenerItem);
    debugText.append("\n"+info);
}
} //end of addListener() method

/**
 * Adds a new listener to specified channel with a specified priority and
 * a given filter object
 * @param listener A channel listener
 * @param filter Listener's filter
 * @param priority Listener priority
 * @param channel_id Channel identification number
 */
public void addListener(ChannelListener listener, ChannelFilter filter,
    int priority, int channel_id)
{
    ChannelListenerItem newListener = new ChannelListenerItem (listener, filter, priority);
    this.addListener(newListener,channel_id);
} //end of addListener() method

/**
 * Adds a new listener to specified channel with a specified priority
 * @param listener A channel listener
 * @param priority Listener's priority
 * @param channel_id Channel identification number
 */
public void addListener(ChannelListener listener,
    int priority,int channel_id)

```

```

{
    ChannelListenerItem newListener = new ChannelListenerItem (listener, priority);
    this.addListener(newListener, channel_id);
} //end of addListener() method

/**
 * Adds a new listener to specified channel with a given filter object
 * @param listener A channel listener
 * @param filter Listener's filter
 * @param channel_id Channel identification number
 */
public void addListener(ChannelListener listener, ChannelFilter filter,
                        int channel_id)
{
    ChannelListenerItem newListener = new ChannelListenerItem (listener, filter);
    this.addListener(newListener, channel_id);
} //end of addListener() method

/**
 * Adds a new listener to specified channel without priority and filter
 * @param listener A channel listener
 * @param channel_id Channel identification number
 */
public void addListener(ChannelListener listener, int channel_id)
{
    ChannelListenerItem newListener = new ChannelListenerItem (listener);
    this.addListener(newListener, channel_id);
} //end of addListener() method

/**
 * Removes a registered listener from specified channel
 * @param listener listener object to be removed
 * @param channel_id Channel identification number
 */
public void removeListener(ChannelListener listener, int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel)memberChannels.get(key_id)).removeListener(listener);
    }
    else
    {
        System.out.println("There is no channel with this id for removing the listener" +
            "\n Unfound channel id: "+ channel_id);
    }
} //end of removeListener() method

/**
 * Removes a registered talker from specified channel
 * @param talker Talker object to be removed

```

```

* @param channel_id Channel identification number
*/
public void removeTalker(Object talker, int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel)memberChannels.get(key_id)).removeTalker(talker);
    }
    else
    {
        System.out.println("There is no channel with this id for removing the talker" +
            "\n Unfound channel id: " + channel_id);
    }
}
} //end of removeTalker() method

/**
* Add the specified filter object to desired channel listener.
* @param listener Channel listener
* @param newFilter Filter object,which will be added to listener
* @param channel_id Channel identification number
*/
public void addListenerFilter(ChannelListener listener, ChannelFilter filter, int
channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel)memberChannels.get(key_id)).addFilter(listener,filter);
    }
    else
    {
        System.out.println("There is no channel with this id for adding a listener filter" +
            "\n Unfound channel id: " + channel_id);
    }
}
} //end of addListenerFilter() method

/**
* Remove the specified filter object from desired channel listener.
* @param listener Channel listener
* @param oldFilter Filter object,which will be removed from the listener
* @param channel_id Channel identification number
*/
public void removeListenerFilter(ChannelListener listener, ChannelFilter filter, int
channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel)memberChannels.get(key_id)).removeFilter(listener,filter);
    }
}

```

```

else
{
    System.out.println("There is no channel with this id for removing a listener filter" +
        "\n Unfound channel id:"+ channel_id);
}
} //end of removeListenerFilter() method

/**
 * Remove all filter objects from desired channel listener.
 * @param listener Channel listener
 * @param channel_id Channel identification number
 */
public void removeAllListenerFilters(ChannelListener listener,int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) ){
        ((Channel)memberChannels.get(key_id)).removeAllFilters(listener);
    }
    else
    {
        System.out.println("There is no channel with this id for removing all listener filters" +
            "\n Unfound channel id:"+ channel_id);
    }
} //end of removeAllListenerFilters() method

/**
 * Starts self-dispatching for a listener on specified channel
 * @param listener Channel listener to be upgraded to self dispatching
 * @param channel_id Channel identification number
 */
public void startListenerSelfDispatch(ChannelListener listener, int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel)memberChannels.get(key_id)).startListenerSelfDispatch(listener);
    }
    else
    {
        System.out.println("There is no channel with this id for starting self-dispatching" +
            "\n Unfound channel id:"+ channel_id);
    }
} //end of startListenerSelfDispatch() method

/**
 * Suspends self-dispatching for a listener on specified channel.
 * @param listener Channel listener to be downgraded to channel dispatching
 * by suspending
 * @param channel_id Channel identification number
 */
public void suspendListenerSelfDispatch(ChannelListener listener, int channel_id)

```

```

{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel)memberChannels.get(key_id)).suspendListenerSelfDispatch(listener);
    }
    else
    {
        System.out.println("There is no channel with this id for suspending self-dispatching"
+
                "\n Unfound channel id:" + channel_id);
    }
}
//end of suspendListenerSelfDispatch() method

/**
 * Resumes a suspended self-dispatcher for a listener on specified channel.
 * @param listener Channel listener to be upgraded to self-dispatching
 * by resuming its existing suspended self-dispatcher
 * @param channel_id Channel identification number
 */
public void resumeListenerSelfDispatch(ChannelListener listener, int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) ){
        ((Channel)memberChannels.get(key_id)).resumeListenerSelfDispatch(listener);
    }
    else{
        System.out.println("There is no channel with this id for resuming self-dispatching" +
                "\n Unfound channel id:" + channel_id);
    }
}
//end of resumeListenerSelfDispatch() method

/**
 * Stops self-dispatching for a listener on specified channel.
 * This process will end the self-dispatching thread.
 * @param listener Channel listener to be downgraded to channel dispatching
 * by stopping self-dispatching
 * @param channel_id Channel identification number
 */
public void stopListenerSelfDispatch(ChannelListener listener, int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        ((Channel)memberChannels.get(key_id)).stopListenerSelfDispatch(listener);
    }
    else
    {
        System.out.println("There is no channel with this id for stopping self-dispatching" +
                "\n Unfound channel id:" + channel_id);
    }
}

```

```

} //end of stopListenerSelfDispatch() method

/**
 * Test if the specified channel has any registered listener
 * @param channel_id Channel identification number
 * @return Returns true if this channel has any registered listener,
 * false otherwise
 */
public boolean hasListeners(int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        return ((Channel)memberChannels.get(key_id)).hasListeners();
    }
    else
    {
        System.out.println("There is no channel with this id " +
            "\n Unfound channel id:" + channel_id);
        return false;
    }
} //end of hasListeners() method

/**
 * Test if the specified channel has any registered talker
 * @param channel_id Channel identification number
 * @return Returns true if this channel has any registered talker,
 * false otherwise
 */
public boolean hasTalkers(int channel_id)
{
    Integer key_id = new Integer(channel_id);
    if(memberChannels.containsKey(key_id) )
    {
        return ((Channel)memberChannels.get(key_id)).hasTalkers();
    }
    else
    {
        System.out.println("There is no channel with this id " +
            "\n Unfound channel id:" + channel_id);
        return false;
    }
} //end of hasTalkers() method

/**
 * Gets the member channels.
 * @return Returns member channels.
 */
public Hashtable getMemberChannels()
{
    return (Hashtable) memberChannels.clone();
}

```

```

} //end of getMemberChannels() method

/**
 * Test if this channel manager has an authority.
 * @return Returns true if this channel manager has an authority,
 * false otherwise
 */
public boolean hasAuthority()
{
    return hasAuthority;
} //end of hasAuthority() method

/**
 * Gets channel manager debug statements
 * @return Returns channel manager debug statements
 */
public JTextArea getManagerDebugText()
{
    return debugText;
} //end of getManagerDebugText() method

/**
 * Removes a listener from all channels.
 * @param listener Listener object to be removed
 */
public void removeListenerFromAllChannel(ChannelListener listener)
{
    Enumeration e = memberChannels.elements();
    while(e.hasMoreElements())
    {
        Channel ch = (Channel) e.nextElement();
        if(ch.isRegisteredListener(listener) )
        {
            ch.removeListener(listener);
        }
    }
} //end of removeListenerFromAllChannel() method

/**
 * Removes a talker from all channels.
 * @param talker Talker object to be removed
 */
public void removeTalkerFromAllChannel(Object talker)
{
    Enumeration e = memberChannels.elements();
    while(e.hasMoreElements())
    {
        Channel ch = (Channel) e.nextElement();
        if(ch.isRegisteredTalker(talker) )
        {
            ch.removeTalker(talker);
        }
    }
} //end of removeTalkerFromAllChannel() method

```

```
    }  
  }  
  } //end of removeTalkerFromAllChannel() method  
} //end of ChannelManager
```


THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX B. THE ADDITIONAL CLASS FOR SAAM PROTOTYPE

```
package org.saamnet.saam.util;

/**
 * The PermissionTableEntry class was created to determine whether or not
 * a channel participant (a talker or listener) has access to a given
 * channel. A channel permission entry is created by providing the
 * identification number of the channel and specifying the channel participants
 * allowed to access this channel.
 */

public class PermissionTableEntry
{

    /**Keeps the identification number of the channel
     * associated with this permission entry
     */
    private int channel_id;

    /** Keeps the verified talkers */
    private String [] validTalkers ;

    /** Keeps the verified listeners */
    private String [] validListeners;

    /** Constructs a permission entry with given channel id and
     * verified channel participants
     * @param ch_id Channel id
     * @param talkers Verified talkers
     * @param listeners Verified listeners
     */
    public PermissionTableEntry ( int ch_id,
                                String [] talkers,
                                String [] listeners)
    {

        channel_id = ch_id;
        validTalkers = talkers;
        validListeners = listeners;
    }

    /**
     * To check whether or not a talker is allowed to access this channel
     * @param talker A channel talker to be checked
     * @return Returns true if this talker is verified to access the channel
     * else false.
     */
}
```

```

*/
public boolean isValidTalker(String talker)
{
    boolean result = false;
    for(int i=0;i<validTalkers.length;i++)
    {
        if( validTalkers[i].equals(talker) )
        {
            result = true;
            break;
        }
    }
    return result;
}

/**
 * To check whether or not a listener is allowed to access this channel
 * @param listener A channel listener to be checked
 * @return Returns true if this listener is verified to access the channel
 * else false.
 */
public boolean isValidListener(String listener){
    boolean result = false;
    for(int i=0;i<validListeners.length;i++)
    {
        if( validListeners[i].equals(listener) )
        {
            result = true;
            break;
        }
    }
    return result;
}

/**
 * This method returns a static permission table (an array of permission table entries)
 * containing the permissions for all channels, which are currently used in SAAM
 * prototype, after the integration.
 * @return Returns current Saam channel permissions
 */
public static PermissionTableEntry [] getSaamPermissions()
{
    PermissionTableEntry [] permissionTable = {

        new PermissionTableEntry(80010,
            new String []{"org.saamnet.saam.control.ControlExecutive"},
            new String []{"org.saamnet.saam.router.TransportInterface"}),

        new PermissionTableEntry(80020,
            new String []{"org.saamnet.saam.router.TransportInterface",
                "org.saamnet.saam.router.Interface"},

```

```

        new String []{"org.saamnet.saam.router.RoutingAlgorithm"}),

new PermissionTableEntry(80030,
    new String []{"org.saamnet.saam.router.RoutingAlgorithm"},
    new String []{"org.saamnet.saam.router.TransportInterface"}),

new PermissionTableEntry(80040,
    new String []{"org.saamnet.saam.router.RoutingAlgorithm"},
    new String []{"org.saamnet.saam.router.Interface"}),

new PermissionTableEntry(80050,
    new String []{"org.saamnet.saam.router.Interface"},
    new String []{"org.saamnet.saam.agent.router.Scheduler"}),

new PermissionTableEntry(80060,
    new String []{"org.saamnet.saam.router.NetworkInterfaceCard",
        "org.saamnet.saam.router.RoutingAlgorithm"},
    new String []{"org.saamnet.saam.router.Interface"}),

new PermissionTableEntry(80070,
    new String []{"org.saamnet.saam.agent.router.Scheduler"},
    new String []{"org.saamnet.saam.router.NetworkInterfaceCard"}),

new PermissionTableEntry(80080,
    new String []{"org.saamnet.saam.router.NetworkInterfaceCard"},
    new String []{"org.saamnet.saam.Translator"}),

new PermissionTableEntry(80090,
    new String []{"org.saamnet.saam.Translator",
        "org.saamnet.saam.Translator$PortListener"},
    new String []{"org.saamnet.saam.router.NetworkInterfaceCard"}),

new PermissionTableEntry(80100,
    new String []{"org.saamnet.saam.Translator",
        "org.saamnet.saam.Translator$PortListener",
        "org.saamnet.saam.router.TransportInterface"},
    new String []{"org.saamnet.saam.control.PacketFactory"}),

};
return permissionTable;
}

/**
 * This method returns the identification number of the channel associated with this
 * entry
 * @return Returns the channel identification number.
 */
public int getValidChannelId()
{

```

```
    return channel_id;
}
} //End of PermissionTableEntry class
```

LIST OF REFERENCES

- [1] Rege, K., “Design patterns for component-oriented software development”, EUROMICRO Conference, 1999. Proceedings. 25th, Volume: 2, 1999.
- [2] Karin Högstedt, Doug Kimelman, V.T. Rajan, Tova Roth, Mark Wegman, Nan Wang, “Optimizing Component Interaction”, ACM SIGPLAN Notices Volume 36 Issue 8, August 2001.
- [3] Vrable, Dean J. and Yarger, John W., “The SAAM architecture: enabling integrated services”, Computer Science Department, Naval Postgraduate School, Monterey, September 1999.
- [4] Grant, Palmer, “Java Event Handling”, Prentice Hall Ptr, 2001.
- [5] Niemeyer, Patrick and Knudsen Jonathan, “Learning Java”, O’reilly, 2000.
- [6] Object Management Group, “Corba Event Service Specification Version 1.0”, June 2000.

[http://www.omg.org/technology/documents/spec_catalog.htm]
- [7] Siegel, Jon, “CORBA 3 Fundamentals and Programming”, John Willey & Sons, 2000.
- [8] Iona, “Orbix Event Programmer’s Guide”, 2001.

[http://www.iona.com/docs/manuals/orbix/33/html/orbixevents33_pgguide/intro.html]
- [9] Bartlett, Dave, “Corba Junction: Corba 3.0 Notification Service”, May 2001.

[<http://www-106.ibm.com/developerworks/components/library/co-cjct8/index.html>]
- [10] Paul, Hyde, “Java Thread Programming”, Sams, 1999.
- [11] Goetz, Brian, “Threading lightly: Synchronization is not the enemy”, July 2001.
[<http://www-106.ibm.com/developerworks/library/j-threads1>]
- [12] Meyer, Bertrand, Object-Oriented Software Construction, Prentice Hall, 1988.

- [13] Berg, Daniel J. and Fritzinger, Steven, “Advanced Techniques for Java Developers, Proven Solutions from Leading Java Experts”, John Wiley & Sons, Inc., 1997.
- [14] Object Management Group, “Corba Notification Service Specification Version 1.0”, June 2000.
[\[http://www.omg.org/technology/documents/spec_catalog.htm\]](http://www.omg.org/technology/documents/spec_catalog.htm)
- [15] Moon, Jae-Chul, Park Jun-Ho, Kang Soon-Ju, “An event channel-based embedded software architecture for developing telemetric and teleoperation systems on the WWW”, Real-Time Technology and Applications Symposium, Proceedings of the Fifth IEEE, 1999.
- [16] Zhou, Dong, Schwan, K., Eisenhauer, G., Chen, Yuan, “JECho - interactive high performance computing with java event channels”, Parallel and Distributed Processing Symposium, Proceedings 15th International, 2001
- [17] Dao-Chang, K., Gibson, John H., “Design of a Dynamic Management Capability for SAAM System to Support Requests For Guaranteed Quality Of Service Traffic Routing and Recovery”, Computer Science Department, Naval Postgraduate School, Monterey, September 2000.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Deniz Kuvvetleri Komutanligi
Personel Daire Baskanligi
Bakanliklar
Ankara, TURKEY
4. Deniz Harp Okulu Komutanligi
Kutuphanesi
Tuzla
Istanbul, TURKEY
5. Chairman, Code CS
Naval Postgraduate School
Monterey, California
6. Prof. Geoffrey Xie, Code CS/Xg
Naval Postgraduate School
Monterey, California
7. LCDR Chris Eagle
Naval Postgraduate School
Monterey, California
8. LTJG Cihat Eryigit
Arastirma Merkezi Komutanligi
Pendik
Istanbul, TURKEY
9. LTJG Tolga Demirtas
Deniz Harp Okulu Komutanligi
Tuzla
Istanbul, TURKEY