

Framework Integrating Real-Time Specific Remote Method Invocation

for Java's

Andy Wellings
University of York
andy@cs.york.ac.uk

Rajkumar Douglens
The MITRE Corporation
{rkjensen}@mitre.org

Doug Wells
The Open Group
dmw-java@contek.com

Abstract

The Distributed Real-Time Specification for Java (DRTSJ) is being developed in the Sun's Java Community Process. It focuses on supporting predictable end-to-end timeliness for sequentially distributed computations. This paper reports an investigation that extends the existing Real-Time Specification for Java and Java Remote Method Invocation facilities to provide the facilities for

remote procedure call (RPC) and remote method invocation (RMI)

- *dataflow*: movement of data without execution point among application entities, e.g., publish/subscribe data transfers
- *networked*: asynchronous movement of messages without execution point among application entities, e.g., message passing in PC

Other distributed system programming models, e.g., mobile objects, autonomous agents, tuple spaces, web services, etc., are currently confined to local or networked real-time systems. For example, most of the approaches have provided augmented distributed models in Java (for example, [25]), JavaSpace [24], Voyage [20], and JavaPart [17]) have not considered real-time issues [1].

Control models are usually designed for multi-node systems, usually *trans-node* (linearly sequential) behavior that is asynchronous, where requests are executed in a first-come, first-served order. When asynchronous responses are needed, one-way invocation can be used to spawn asynchronous tasks. An example of this distributed control approach is the distributed thread model. A *distributed thread* is a single thread with a system-wide ID, that extends and retracts itself sequentially through an arbitrary number of local and remote objects. A *distributed real-time thread* transparently propagates its timeliness properties (and perhaps resource ownership, transactional context, security attributes, etc.) when its execution point transits object and perhaps node boundaries. The distributed thread model appeared in the Alpha distributed real-time OS kernel [1] and was subsequently incorporated in the JDK 7 distributed real-time [26]; it is the basis of the programming model for recent Real-Time CORBA (Dynamic Scheduling) specification [14]. Realistic non-trivial experimental distributed real-time computing systems are successfully constructed in Alpha [26] and [7].

Almost all of the models including those for real-time 'publish/subscribe' are more oriented toward maximizing throughput than maintaining end-to-end timeliness properties (cf. OMG's RFP for Data Distribution Service Real-Time System [15]). A real-

1 Introduction

The Distributed Real-Time Specification for Java (DRTSJ) is being developed by SR-50 experts in the Sun Java Community Process [7]. One of the issues being explored is the options available for integrating the Real-Time Specification for Java (RTSJ) [2] and Java Remote Method Invocation (RMI) facility [23].

This paper is structured as follows. Section 1 motivates the need for integrating RTSJ and RMI. Section 2 discusses the points of integration. Section 3 proposes a minimal set of changes to integrate RTSJ with RMI. The limitations of the proposal are detailed leading to the introduction of section 4, the notion of Real-Time Remote Interface along with the discussion of this proposal. Section 5 discusses the proposal and evaluates it. Section 6 discusses the proposal and discusses distributed thread functionality and interoperability between the various platforms. Finally, section 8 summarizes the proposals and presents conclusions for the future work.

2 Distributed Real-Time Systems

There are many different kinds of distributed systems. However, the majority of deployed distributed time computing systems employ one of the following programming models:

- *control flow*: movement of execution point with parameters among application entities, e.g.,

time flow specific information outside the scope of RTSJ. The message passing provided by typical real-time distributed real-time systems forces end-to-end timeliness and integrity responsibility. Most frequently distributed real-time become the responsibility of each application programmer. Frequently custom distributed time middleware is created. As commercial distributed real-time infrastructure products (notably real-time CORBA-compliant ones) continue to emerge, they are becoming widely used.

Java already includes a model for distributed Remote Method Invocation (RMI) [23] and SR-50 proposed an enhancement for real-time systems RMI provides a familiar distributed object system on flow model of method invocations using abstract interface definitions for remote objects and distributed object system programming. A similar asynchronous message passing is provided by allowing synchronous (one-way with return parameters) method invocations. RMI also transfers object instances by value. This allows messages to pass as objects. This supports simple point-to-point flow of objects. It does not effectively publish/subscribe models.

But RMI intentionally provides very limited support for end-to-end properties in particular non-timed real-time systems. Therefore the has the objective of integrating RMI with RTSJ support for end-to-end timeliness by creating control models for programming models. It is a possible mechanism for the distributed system model such as point-to-point data flow messaging and mobile objects.

A distributed object system (or application services) requires a well-defined architecture stable interface between components and a system-wide agreement on infrastructure technology. real-time distributed object system that agreement includes the semantics of timeliness and sufficient synchronization to avoid deadlocks. Typically such systems found in distributed enterprise are difficult to build technical agreement and coordination needed to build distributed object system between enterprises. This is consistent with the current normal deployment of distributed real-time computing systems. The callab requirements of RTSJ based on the presumption of intra-enterprise environments.

3 RMI in Java

This paper assumes that the reader is familiar with RTSJ and RMI. However, in order to integrate the necessary interpretation of the design goals

current OS used any high level time on all ed e , ms control tion ided er be data that ed awas DRTSJ to flow els va's tance, with of kn ly are at the ad is ility ion the at, abs

RMI purposes, the main component of RMI can be considered as follows:

- the programming model where the objects that are accessed remotely are identified via a remote interface,
- the implementation model, which provides the transport mechanisms where by a Java platform can talk to another node to request access to objects and
- the development tools (e.g. rmic) its dynamic counterpart which takes remote objects and generate the proxies required to facilitate communication.

Key to developing RMI-based systems is defining the interface to the remote object. RMI requires that the programmer extend the pre-defined interface. RemoteEach method defined in the interface extending RemoteEach throws RemoteException. Thus on the design decision of RMI is that it is not completely transparent to the programmer. The location of the remote objects is transparent, but the more necessary is transparent.

3.1 The meaning of time

Standard Java times are expressed as a number of milliseconds and nanoseconds as a midnight. In Java 1.970. The types of these values are long and int respectively. java.util.Date class encapsulates these values. The Date class is serializable and consequently objects of this type can pass through RMI.

RMI is not a simple relationship between the client and the server site. However, the leasing mechanism of the distributed garbage collection algorithm is similar to the lock progress approximation algorithm. If this assumption is violated, remote objects might unexpectedly disappear. Resulting appropriate remote exceptions are generated when attempted.

3.2 Failure semantics

RMI is a simple reliable transport mechanism (e.g., TCP). Consequently, it does not need to be concerned with transient communication errors. RMI handling of errors and permanent communication failures are centered around the throwing of a RemoteException. A RemoteException is thrown if a remote method call fails. The RMI implementation is then able to make the call detectable as a failure for the returned.

The server can inform the client of details of the server executing request. The RMI implementation phases terminate at 16. Hence RMI is

exactly once semantics in the presence of failures and
 once semantics in the presence of failures [18].

4 Minimal Integration between RTS and RMI Level 1 Integration)

Interface Java provides mechanisms for refining
 contradictions between client and server. They make
 statements about the attributes of objects associated
 [20]. By implementing an interface, the server is
 guaranteeing to provide the functionality implied
 interface. *By definition, a Java interface says nothing
 about the real-time properties of the server or
 other non-functional requirements. Similarly, an
 interface says nothing about the real-time properties
 of the server or the underlying transport system.*
 Consequently, an object that implements a remote
 interface is an object that is executing on a
 real-time platform. Furthermore, the connection
 between the client and server are required to
 be timely, irrespective of whether the client is a
 client. These assumptions are accepted by the
 that generate the client and server proxies and do
 without knowledge of the server or the
 execution platform.

One of the interesting interpretations of RMI
 is that the proxy thread is the thread which executes
 server methods on behalf of the client. In an
 ordinary Java thread, even the client's
 thread proxy thread is viewed as a normal Java
 thread. This is the minimal integration
 between RTS and RMI (called Level 1 integration).
 The threads are remote methods, but they expect
 timely delivery of RMI requests to the server
 proxy thread, and a real-time constraint on the
 has. Consequently, the application programmer must
 explicitly pass any scheduling or release parameter
 between the client and the server and require a
 sympathetic RMI implementation.

The main advantage of Level 1 integration between
 RTS and RMI is that it requires no additional RMI
 threads. Given this, Level 1 integration is a silent
 relationship between the client and server. It also
 has a failure semantics of RMI.

Of course, this is one interpretation of
 using a real-time environment. Another
 implementation mechanism should be defined that
 what the real-time environment does
 the real-time VM. This approach proposes
 because it requires changing RMI. It likely
 time overhead. Moreover, assuming that the
 server attributes are already expressed in the
 difficult to determine the server's

how RMI should be implemented and
 the RMI should identify the server to exploit
 this paper. It has a great advantage of explicit
 client might be more expected.

5 Real-Time RMI Level 1 Integration)

This discussion presents a design that suggests that an
 real-time RTS should make RMI threads
 timely delivery of requests to the server
 with inheritance of the timing requirements of
 the client.

Keeping with the non-transparent RMI philosophy,
 to obtain real-time remote communication Level 1
 integration proposes the introduction of a real-time
 RMI. Real-time RMI consists of the following components.

- the programming model where real-time objects that
 can access remote objects are identified as real-time
 remote interface (this is similar to the approach
 that Miguel has adopted for the addition of
 predictable RMI using reservation-based scheduling
 techniques [10]),
- the implementation model which provides timely
 transport mechanisms whereby the RTS platform
 can allow the order of request access to the
 objects and pass across timing constraints or
 scheduling parameters associated with the client
 and
- the development of (e.g. modified) which
 takes server objects and generate the real-time proxy
 required to facilitate real-time communication.

Key to developing real-time RMI-based systems is
 defining the interface to real-time objects (objects
 that assume they are executing on the platform).
 Real-time RMI requires that objects that provide
 remote interface must indicate by extending the
 defined interface `RealtimeRemote`.

```
public interface RealtimeRemote
  extends java.rmi.Remote {};
```

Each method defined explicitly or implicitly in an
 interface extending `RealtimeRemote` must declare that
 "it throws `RemoteException`".

Also, the proxy thread is viewed as an
 RTSJ `RealtimeThread` that inherits appropriate
 scheduling and release parameters from the client's
 RTSJ. When the client's RTSJ thread is a
 thread, default release scheduling parameters
 provided. Clients that are asynchronous servers and
 need further consideration. Asynchronous servers
 are expected to be lightweight objects. If they are
 objects, this may well add to the cost of their
 implementation. It is assumed that they
 and the proxy thread are created on the server's
 real-time thread.

It is beyond the scope of this paper to address
 locating the real-time server versus associated
 determining the real-time properties of the connection
 possible real-time mechanisms with these
 and JavaSpaces [24] technologies. su associated with
 how the client can
 server. up to [25]

A RealtimeRemote interface indicates that the client expects the underlying transport messages and server-side objects to be aware of any client-side scheduling parameters. However, if a guarantee of the type of memory used is strongly guaranteed, giving the interface a definition like the NoHeapRealtimeRemote interface. This ensures that the underlying transport messages and server objects will have no heap. The server proxy thread can be viewed as an RTSJ NoHeapRealtimeThread.

One issue that must be considered is whether there needs to be a change in the serialization mechanism to support passing objects across time from a no-heap real-time to a heap interface. This is achieved by producing extensions to java.io.ObjectOutputStream and io.ObjectInputStream classes.

5.1 Model

The level of integration of RTSJ and RMI assumes that the communication timing constraint between client and server is transparent to both the client and server RTSJ platforms. For example, a server RTSJ platform is concerned with the proxy thread just as another real-time thread whose scheduling and parameters just happen to be the real-time infrastructure. The client and server RTSJ platforms are therefore, still independent of one and other. Consequently, the relationship between the two locks.

It should be noted that currently RTSJ classes are not serializable. This means that any message passed across the RMI network is converted to a byte array and constructed at the server site.

5.2 Failure semantics

Level of integration assumes that a sufficient crash failure [6].

In the absence of failure, level of integration provides exactly the same semantics. In the presence of failures, the semantics are the same. Failure on the server is presented to the client as a remote exception. Failure on the client is ignored on the server side. Server errors are reported by the RTSJ synchronization mechanisms. Either an asynchronously interrupted exception (AIE) is thrown to the server proxy thread, or a synchronous server (A) error is reported.

5.3 Limitations of Real-Time RMI Level Approach

Although the level of integration requires extension of RMI, it does not require any extension to the RTSJ.

the RTSJ. However, this results in a set of limitations of the approach.

- limitation of the RMI definition with the interaction
 - limitation of the RTSJ definition
- The limitation from RMI/JVM that is a concern here is the asynchronous nature of RMI. The following situation will result in a deadlock: a call to a synchronized method of an object which all remote methods of that object call. The object is directly or indirectly synchronized to the JVM. In RMI, the object is a collective of distributed threads, each executing the method of the object. A thread is considered a single thread. Not that a single thread is performed on a per-thread basis. There would be a deadlock if the object is a single thread and is only local to the remote method invocations.

A similar problem occurs with thread-local data. However, the problem is not as severe as the case with the object. In the case of the original object, the thread is the proxy thread and is not considered the original thread.

One way to solve the problem is to require the introduction of a global thread identifier, as carried with the call. The JVM and RTSJ need to be changed in order to check this identifier. For the synchronized method situation described above, the RTSJ VM must allow threads to be synchronized. This raises a remote exception indicating that a deadlock occurred. Allowing the proxy thread to enter while conceptually simple has a performance cost. Essentially, under certain failure conditions, implementation will result in a thread being a synchronized object. Raising an exception is a simple implementation of a global thread identifier. Thread-local data is not correctly.

The limitation imposed by the RTSJ is more difficult to circumvent. The RTSJ class definitions define remote interfaces. Consequently, they can offer remote services. Furthermore, with the exception of Asynchronously Interrupted Exception, none of the classes implement the Serializable interface, and consequently objects of these classes cannot be passed across a remote interface. The next section considers extensions to the RTSJ to increase its distributed real-time programming capabilities.

6 Distributed Real-Time Thread Level Integration

The distributed real-time thread model is a section of which thread local control can move freely across distributed systems by a method of remote objects. Each distributed deal-time thread has a unique system-wide identifier and a point in time (in the absence of faults) such as network partitioning, thread migration, or is suspended in a single distributed system. This is called the **head** of the thread. It is the distributed thread object that is created at the **root origin** of the thread, the hosting of the distributed thread is called **segment**. Any remote operation of the distributed thread is performed via the underlying real-time virtual machines) on a site that currently has the distributed thread's execution.

One of pressing distributed real-time threads is extended RTSJ's `RealtimeThread` class to implement real-time monitor interface which defines remote operations that can be called on the thread. However, to implement distributed threads requires from the underlying RTA platform a real-time transport protocol that can implement the discussion section. Consequently, in this section of distributed real-time JVM (DRTJVM) and distributed real-time RMI introduced Distributed real-time consist of the following components:

- the programming model where distributed real-time objects that are necessary to identify a `DistributedRealtimeRemote` interface;
- the implementation model which provides timely transport mechanisms where beyond RTSJ platform call in the request objects and pass scheduling parameters and information associated with the identifier to facilitate implementation of distributed thread model;
- the development tool (e.g. a modified real-time rmic) which serve objects generate the real-time time proxies required to facilitate the real-time communication, the serve proxy thread can be considered distributed real-time thread.

The distributed real-time Java platform is a platform augmented with facilities to support the distributed real-time thread model.

6.1 The Distributed Real-Time Specification

Although beyond the scope of this paper, a complete distributed real-time specification of (DRTSJ) is following discussed the main component associated with the distributed thread model.

Distributed Real-Time Threads

The operations that are performed on distributed real-time threads are classified into two areas:

- *Operations that affect the scheduling of the distributed thread.* These include being able to get and set scheduling parameters. One of the motivations for providing these remote operations comes from an assumption that threads in a distributed system might need to perform mode change threads whose origins are other nodes or threads which need to manipulate the threads release and scheduling parameter. (Of course, the functionality you achieved by having the application define new objects to call the threads origins and those objects to perform the mode change. However, allowing direct mode changes offers more flexibility. For example, it allows the underlying DRTSJ platform to optimize the location of the thread to carry the operation. It also allows of tolerance whether the threads are accessible due to network or communication failures.)

- *Operations that affect the execution state of the distributed thread.* These include being able to remotely interrupt. While this is a clear requirement from a distributed thread another mode of a thread that being able to interrupt a thread should be more operation.

This is the focus of the discussion in the following. Figure 1 illustrates the remote operations that are defined on the `DistributedRealtimeThread` class. The RTSJ class definition for `ReleaseParameters`, `SchedulingParameters` etc will need to be re-defined (or subclasses created to implement monitor interface classes created. (Note that the reference semantic of a variable in RTSJ if for example the `ReleaseParameters` of a real-time thread are changed then the scheduled implementation of the change is immediately. However, in RMI the parameters are serialized then passed to the other side of the network. Consequently, the interface shows the parameters are remote interfaces to ensure consistency with RTSJ. A new method allow the current remote thread to be found. A similar definition can be derived for `DistributedNoHeapRealtimeThread`.

Asynchronous Events

As well as supporting real-time threads, the RTSJ also supports the notion of events and event handlers. Event handlers are permanently bound to threads in a distributed real-time RTSJ. It is important to note that these are not handlers to be bound to a scheduled thread.

```

public interface RemoteThread extends DistributedRealtimeRemote
{
    RemoteReleaseParameters getReleaseParameters() throws RemoteException;
    void setReleaseParameters(RemoteReleaseParameters parameters)
        throws RemoteException;
    /* Similarly for SchedulingParameters */

    RemoteScheduler getScheduler() throws RemoteException;

    synchronized void interrupt() throws RemoteException;

    void start() throws RemoteException, IllegalThreadStateException;
}

public DistributedRealtimeThread extends RealtimeThread implements RemoteThread
{
    // Implementation of RemoteThread interface plus

    public static RemoteThread currentRemoteThread()
        throws NotARemoteThreadException;
}

```

Figure 10. Distributed Real-time Threads

There are several issues to consider when adapting models for execution in a distributed environment. It is assumed that the same restrictions that were applied to distributed real-time threads also apply to distributed real-time threads. However, distributed real-time threads are not serializable and hence their handlers cannot be persistent. Both event and remote handlers can have remote interfaces and therefore can be invoked on a node as a result of the scheduling handler. Another distributed handler is a distributed real-time thread that can make remote calls.

One possible definition of the interface and class needed for distributed event and remote handlers is shown in Figure 10. The definition is similar to the one for event and remote handlers that attach and manipulate the bind to method to allow remote operation. The only available operation in the AsynchronousEventHandler interface is the remote schedule. The remote schedule method determines which scheduler should be formed when an event has been fired. There is no need for a Runnable interface in the RTSJ hierarchy for event handlers in a distributed environment.

Asynchronous Transfer of Control

The facilities for handling asynchronous transfers of control (ATC) in a language are intrinsically complex.

this is a thread event and land em on buted es given ingh related. Remote- to n dlers.

The facilities for raising an asynchronous exception (using the AsynchronouslyInterruptedException(AIE) class and the Interruptible interface) are not backward compatible with standard Java approaches with an interrupt mechanism. Although ATCs are intrinsically complex, they do not require the same requirements as time-based needs for real-time synchronous event quickly and safely. This underlying need is independent of the threads distributed thread real-time thread.

Unfortunately applying ATCs to distributed threads adds a heretofore complexity to a single process mode. This is because the RTSJ is a many-to-many architecture that accesses threads directly (moving the address of the thread object being executed) rather than through a directed (moving the address of the thread) may already be outstanding AIEs of the thread. Careful consideration needs to be given to an implementation model that will support the current semantics efficiently. Either,

1. outstanding AIEs will be able to interrupt the thread and become interruptible. This check will be done by throwing this out to be prohibitive expensive unless further restrictions are made. For example, if a thread is not caught, further more the approach would be a failure or,

if event serializable problem would occur with event as interrupt current could be a single thread would have passed to the network handler could be a single thread serializable however the model would be confused and not being scheduled.

```

public interface RemoteAsynchronousEvent extends
    DistributedRealtimeRemote
{
    void addRemoteHandler( RemoteAsynchronousEventHandler handler)
        throws RemoteException;
    // also handledBy, removeHandler, setHandler
    // createReleaseParameters for remote handlers

    void fire() throws RemoteException;
}

public class GlobalAsyncEvent extends AsyncEvent
    implements RemoteAsynchronousEvent { ... };

public interface RemoteAsynchronousEventHandler
    extends DistributedRealtimeRemote
{
    RemoteScheduler getScheduler() throws RemoteException;
}

public class GlobalAsyncEventHandler extends AsyncEventHandler
    implements RemoteAsynchronousEventHandler;

public class GlobalBoundedAsyncEventHandler
    extends GlobalAsyncEventHandler;

```

Figure 2: Asynchronous Events

2. outstanding AIEs carried the
offload⁴.

Itattens find in the adiffi
least expensive, current the facilit
Assuming that AIEs distributed threads requ
necessary to decide the remote operation tha
performed then the operation on AIEs be
partitioned into sets that require the
threads and those required by the inter
Currently, the TSOs are not making this distinction
consequently, threads firing AIEs also call
methods to enable/disable the AIEs in those range
happen in the method⁵.

One possibility is to define two interfaces.
first for threads that wish to fire the AIE
second interface for those off threads that will
interrupted. These interfaces along with the other
associated classes are given in figure 2. Here,
threads will not execute the doInterruptible
method - threads therefore only execute the

ad
but
used.
ired,
cthe
interrupting
rupted.
clear;
he
the
The
The
be

doInterruptibleThe method however can
make remote as subsequent parameters
must be changed to take an interface parameter A
restriction that the interruptAction should be local
allow the handling of time access parameter
might be declared in the IE.

Given that Asynchronously Interrupted Excep-
tion implements Serializable, consideration
needs to be given as to whether it is appropriate for
distributed asynchronous interrupted exception to
pass across the network. This is allowed, what
semantics doInterruptible progress?

Having interrupted the remote interface of a
DistributedRealtimeThread allows remote threads
to be generic. It is shown.

Exporting Distributed Real-Time Objects

If a reference to remote objects is passed
the network, the objects are exported. This
can be achieved by defining the remote objec
inher from one of the subclasses of
java.rmi.DistributedRealtimeRemoteObject or by
calling the static method exportObject with the
remote object as parameter. This paper has
assumed that the relevant classes will have constructors
that all exportObject.

⁴ Other thread control information can be
the thread, for example security information, whether
daemon, etc.

⁵ The argument would be that asynchronous
possibilities to use remote interface
threads for asynchronous notification without
whether the result is scheduled or not.

⁵ ssed the af
h threads

vent and ne
This would allow
being concerned
handler.

```

public interface RemoteFirer extends DistributedRealtimeRemote
{
    boolean fire() throws RemoteException;
}

public interface RemoteAsynchronouslyInterruptedException
    extends DistributedRealtimeRemote
{
    boolean disable() throws RemoteException;
    boolean enable() throws RemoteException;
    boolean isEnabled() throws RemoteException;
    boolean happened(boolean propagate) throws RemoteException;
}

public class DistributedAsynchronouslyInterruptedException
    extends AsynchronouslyInterruptedException
    implements RemoteFire, RemoteAsynchronouslyInterruptedException
{
    // implement the interfaces and
    public boolean doInterruptible(DistributedInterruptible logic);
}

public interface DistributedInterruptible
{
    void interruptAction (AsynchronouslyInterruptedException exception);
    void run (RemoteAsynchronouslyInterruptedException exception)
        throws AsynchronouslyInterruptedException;
}

```

Figure 3: Asynchronous Transfers of Control

6.2 The meaning

Level integration assumes that the cluster of the distributed real-time application should therefore provide perhaps from the real-time lock that is coordinated across the cluster, e.g. synchronized to some defined accuracy (FTC).

tes
of the
parate
[13]
and

- When the distributed thread returns to the failed segment, the remote exception raised in the segment is previous to the failed site.
- If failed in the origin site approach followed in the distributed thread attempts to return to the origin site, the remote exception is.

6.3 Failure semantics

Sites assumed to differ on a failure on the distributed thread on the DRTS platform directly supporting the distributed thread semantic. Consequently when a segment in the distributed thread fails, the DRTS platform coordinates their response as follows.

y.
is
s.
ibuted

When the thread handling the remote exception is running, a handler details the failure found on the underlying DRTS platform.

AIE
abe

7 Interoperability

The following table shows the level of interoperability between the various models. Note that for each thread model, the subject to be considered is the time of the distributed real-time model. The no-heaviness of the introduction of the requirements for the page.

lity
other
ATC.
tion
real-
aces.
her

- If failed in the origin site, the previous segment has a remote exception raised.
- If failed in the segment, the origin thread implements the model proposed for the site level integration (i.e. ignore through AIE)

Client Class	Remote Interface	Real-time Remote	Distributed Real-time Remote
Thread	Defined by <code>RemoteInterface</code> Server proxy thread view standard Java thread Not Expressible from client	Default timing properties used by real-time <code>Remote</code> Server proxy thread view standard Java thread with default scheduling parameters Not Expressible from client	Default timing properties used by distributed real-time <code>Remote</code> Server proxy thread view standard Java thread with default scheduling parameters Not Expressible from client
Real-Time Thread (Bound) Asynchronous Event Handler	Defined by <code>RemoteInterface</code> Server proxy thread view standard Java thread Not propagated through <code>Remote</code> method invocation in the distributed thread model	Client timing properties used by real-time <code>Remote</code> Server proxy thread view standard Java thread with inherited scheduling parameters and sporadic release parameters All static <code>Remote</code> methods however, in distributed thread model, only client thread interrupted this propagated through real-time <code>Remote</code> system proxy thread	Client timing properties used by distributed real-time <code>Remote</code> Server proxy thread view standard Java thread with appropriate release scheduling parameters All static <code>Remote</code> methods in client thread interrupted this propagated through distributed thread in the DRTSJ platform (the client distributed thread)
Distributed Real-Time Thread (Bound) Global Event Handler	Defined by <code>RemoteInterface</code> Server proxy thread view standard Java thread Not propagated through <code>Remote</code> method invocation in the distributed thread model	Client timing properties used by real-time <code>Remote</code> Server proxy thread view standard Java thread with inherited scheduling parameters All static <code>Remote</code> methods however, in distributed thread model only client distributed thread interrupted this propagated through real-time <code>Remote</code> system proxy thread	Client timing properties used by distributed real-time <code>Remote</code> Server proxy thread view standard Java thread segmented in distributed real-time Java thread All static <code>Remote</code> methods in distributed thread interrupted this has access to distributed real-time objects associated with this propagated through DRTSJ platform (the client distributed thread)

8 Conclusions and Future Work

This paper explores the way in which the DRTSJ can be integrated into the JVM incremental approach. It has been suggested along the following lines in order to increase functionality (and increase complexity).

- Real-time Java threads can be remote objects but they expect timely message delivery and inheritance of scheduling parameters and change of scheduling parameters.

can
has
of
and
there

- Real-time Java threads can call real-time remote objects and they can expect timely delivery of messages and inheritance of scheduling parameters; however they cannot expect to have any distributed thread functionality – the extensions required for these extensions require DRTSJ in the JVM.
- Real-time Java threads can call real-time remote objects and they can expect timely delivery of messages, inheritance of scheduling parameters and

an

associate global thread identifier. However they have not the distributed real-time functionality are extensions required. RMI but extensions required. RTS but small extensions the supporting platform required.

- Distributed real-time threads and distributed real-time objects to the expected timely delivery of messages, inheritance of scheduling parameters, distributed real-time functionality. RTS basis of extensions required. RMI supporting platform.

Currently authors are considering the details Real-time infrastructure and the approach distributed real-time scheduling. Expect that will follow closely the approach adopted dynamic time CORBA [4].

Acknowledgement

The search reported in this paper is performed in the context of the Community Process SR-50. The authors gratefully acknowledge the help and advice by members of SR-50 expert group.

References

1. M. Bogdan, *Distributed Systems*, Wiley, 2001.
2. G. Bollella, *The Real-Time Specification for Java*, Addison Wesley, 2000.
3. Burns and Wellings, *Real-Time Systems and Programming Languages 3*, Addison Wesley, 2001.
4. R. Clark, E. Densen and B. Reynolds, *An Architecture Overview of the Real-Time Distributed Kernel*, USENIX Workshop on Microkernel and Kernel Architectures, 20-28 1993.
5. Clark et al, *Adaptive Distributed Airborne Tracking System*, Workshop on Parallel and Distributed Real-Time Systems, IEEE, <http://www.real-time.org/docs/wpdrts99.pdf>, 1999.
6. F. Cristian, *Understanding Fault-Tolerant Distributed Systems*, ACM, 4(2):16-78, 1991.
7. Eensen, *The Distributed Real-Time Specification for Java*, Initial Proposal, *Journal of Computer Systems Science and Engineering*, March 2001.
8. B. Liskov et al, *Orphan Detection*, IEEE Fault-Tolerance Computing Symposium, July 1987.
9. Maynard et al, *Example Real-Time Command, Control and Management Application for Alpha Architecture*, TR-8812, CMU Computer

Science Dept., <http://www.real-time.org/docs/c2demo.pdf>, December 1988.

10. M. Adigue, *Solution of a Java-RMI Predictable Four-Phase International Symposium on Object-Oriented Real-Time Distributed Computing*, pp 379-388, 2001.
11. J. Northcutt, *Mechanisms for Reliable Distributed Real-Time Operating Systems*, The Alpha Kernel, Academic Press, 1987.
12. R. Berg, *Mastering RMI*, Wiley, 2001.
13. OMG, *Enhanced DCE/TCM, Revised Submission*, OMDocument bos/99-10-02, October 1999.
14. OMG, *Dynamic Scheduling Real-Time CORBA0*, Joint Final Submission OMDocument bos/2001-04-01.
15. OMG, *Distributed Service Real-Time System*, Request for Proposal CORBOS/2001-09-11, Object Management Group, September 2001.
16. F. Panzieri, K. Shrivastava, Rajdod, *Remote Procedure Call Mechanism Supporting Orphan Detection and Killing*, IEEE SOSE4 (1):30-37, 1988.
17. M. Philipps and M. Zenger, *Java Party: Transparent Remote Objects in Java*, *Concurrency Practice and Experience* 9(1):125-124, 1997.
18. P. Park, E. McNiff, *The Remote Method Invocation Guide*, Addison Wesley, 2001.
19. Pitt, *Personal Communication*, October 2001.
20. Recursive Software, *Voyager*, <http://www.recursionsw.com/products/voyager>, accessed February 2002.
21. J. Rumbaugh, G. Jacobson and G. Booch, *The Unified Modeling Language Reference Manual*, Addison Wesley, 1999.
22. Sun Microsystems, *Java Object Serialization Specification*, December 1998.
23. Sun Microsystems, *Java Remote Method Invocation Specification*, December 1999.
24. Sun Microsystems, *JavaSpace Service Specification*, Version 1.0, October 2000.
25. Sun Microsystems, *Java Architecture Specification*, Version 1.0, October 2000.
26. D. Wells, *Trusted Scalable Real-Time Operating System*, Dual-Use Technologies and Applications Conference Proceedings, pp 2-27, 1994.
27. A. Wollrath, *Personal Communication*, October 2001.

-there

uted

y —
I,

of the

support
that

real-

mech

given

An

-

Fault-