

JBoss Transactions 4.2.3

Transaction Core Programmers Guide

TX-PG-4/13/07





**1801 Varsity Drive
Raleigh, NC 27606-2072
USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park, NC 27709
USA**

Documentation-Administration

Copyright © 2007 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of the work or derivative of the work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

Red Hat and the Red Hat "Shadow Man" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

All other trademarks referenced herein are the property of their respective owners.

The GPG fingerprint of the security@redhat.com key is:

CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

Contents

About This Guide	5	Last resource commit optimisation	33
What This Guide Contains.....	5	Nested transactions	34
Audience	5	Independent top-level transactions.....	35
Prerequisites.....	5	Transactions within <code>save_state</code> and	
Organization	5	<code>restore_state</code>	36
Documentation Conventions.....	6	Example	36
Additional Documentation.....	6	Garbage collecting objects	37
Contacting Us	7	Transaction timeouts	37
Overview	8	Hints and tips	39
Introduction.....	8	General	39
TxCore – the transaction engine	8	Using transactions in constructors	39
Saving object states.....	9	More on <code>save_state</code> and <code>restore_state</code>	40
The object store.....	9	Direct use of <code>StateManager</code>	41
Recovery and persistence.....	10	Tools.....	42
The life-cycle of a Transactional Object for		Constructing a Transactional Objects for	
Java.....	11	Java application.....	53
The concurrency controller	12	Application construction	53
The transaction protocol engine.....	13	Queue description	53
Example	14	Constructors and destructors	54
The class hierarchy	14	<code>save_state</code> , <code>restore_state</code> and type	55
Using TxCore	17	enqueue/dequeue operations	56
Introduction.....	17	<code>queueSize</code>	57
State management.....	17	<code>inspectValue/setValue</code> operations	57
Object states.....	17	The client.....	58
The object store.....	19	Comments	59
Selecting an object store implementation	20	Configuration options.....	60
<code>StateManager</code>	20	Options	60
Example	25	Object store implementations.....	62
Lock management and concurrency control	26	The <code>ObjectStore</code>	62
Selecting a lock store implementation	27	Persistent object stores	63
<code>LockManager</code>	27	Class definitions	68
Locking policy	29		
Object construction and destruction	30		
General transaction issues	32		
Advanced transaction issues with TxCore....	32		
Checking transactions	32		

Class library	68
LockManager	68
StateManager	69
Input/OutputObjectState	69
Input/OutputBuffer	70
Uid	70
AtomicAction	71
Index	72



About This Guide

What This Guide Contains

The Transaction Core Programmers Guide contains information on how to use JBoss Transactions 4.2.3. This document provides a detailed look at the design and operation of the TxCore transaction engine and the Transactional Objects for Java toolkit. It describes the architecture and the interaction of components within this architecture.

Audience

This guide is most relevant to engineers who are responsible for administering JBoss Transactions 4.2.3 installations. Although this guide is specifically intended for service developers, it will be useful to anyone who would like to gain an understanding of transactions and how they function.

Prerequisites

This guide assumes a basic familiarity with Java service development and object-oriented programming. A fundamental level of understanding in the following areas will also be useful:

- General understanding of the APIs, components, and objects that are present in Java applications.
- A general understanding of the Windows and UNIX operating systems.

Organization

This guide contains the following chapters:

- **Chapter 1, Overview:** this chapter contains a description of the use of the TxCore transaction engine the Transactional Object for Java classes and facilities.
- **Chapter 2, Using TxCore:** gives details on interfaces and classes defined by TxCore and describes how they can be used to construct transactional applications.
- **Chapter 3, General transactions issues:** presents advanced issues with TxCore.
- **Chapter 4, Hints and tips:** illustrates some hints on the way to use TxCore

- **Chapter 5, Tools:** how to use the management tools shipped with TxCore.
- **Chapter 6, Constructing a Transactional Object for Java application:** this chapter describes a detailed implementation of an application to illustrate various mechanisms provided by TxCore.
- **Chapter 7, Configuration options:** shows configurations options of TxCore.

Documentation Conventions

The following conventions are used in this guide:

Convention	Description
<i>Italic</i>	In paragraph text, italic identifies the titles of documents that are being referenced. When used in conjunction with the Code text described below, italics identify a variable that should be replaced by the user with an actual value.
Bold	Emphasizes items of particular importance.
Code	Text that represents programming code.
Function Function	A path to a function or dialog box within an interface. For example, "Select File Open." indicates that you should select the Open function from the File menu.
() and	Parentheses enclose optional items in command syntax. The vertical bar separates syntax items in a list of choices. For example, any of the following three items can be entered in this syntax: <code>persistPolicy (Never OnTimer OnUpdate NoMoreOftenThan)</code>
Note:	A note highlights important supplemental information.
Caution:	A caution highlights procedures or information that is necessary to avoid damage to equipment, damage to software, loss of data, or invalid test results.

Table 1 Formatting Conventions

Additional Documentation

In addition to this guide, the following guides are available in the JBoss Transactions 4.2.3 documentation set:

- JBoss Transactions 4.2.3 *Release Notes*: Provides late-breaking information about JBoss Transactions 4.2.3.
- JBoss Transactions 4.2.3 *Installation Guide*: This guide provides instructions for installing JBoss Transactions 4.2.3.

- JBoss Transactions 4.2.3 *Failure Recovery Guide*: Provides guidance for administering the system.
- JBoss Transactions 4.2.3 *Transactions API Guide*: Provides guidance for administering the system.
- JBoss Transactions 4.2.3 *Web Services Transactions Programmers Guide*: Provides guidance for administering the system.
- JBoss Transactions 4.2.3 *Administration Guide*: Provides guidance for administering the system.

Contacting Us

Questions or comments about JBoss Transactions 4.2.3 should be directed to our support team.

Overview

Introduction

This chapter contains a description of the use of the *TxCore* transaction engine and the *Transactional Objects for Java* classes and facilities. The classes mentioned in this chapter are the key to writing fault-tolerant applications using transactions. Thus, after describing them we shall apply them in the construction of a simple application. The classes to be described in this chapter can be found in the `com.arjuna.ats.txoj` and `com.arjuna.ats.arjuna` packages.

TxCore – the transaction engine

In keeping with the object-oriented view, the mechanisms needed to construct reliable distributed applications are presented to programmers in an object-oriented manner. Some mechanisms need to be inherited, for example, concurrency control and state management; while other mechanisms, such as object storage and transactions, are implemented as *TxCore* objects that are created and manipulated like any other object.

Note: When the manual talks about using persistence and concurrency control facilities it assumes that the *Transactional Objects for Java (TXOJ)* classes are being used. If this is not the case then the programmer is responsible for all of these issues.

TxCore exploits object-oriented techniques to present programmers with a toolkit of Java classes from which application classes can inherit to obtain desired properties, such as persistence and concurrency control. These classes form a hierarchy, part of which is shown below and which will be described later in this document.

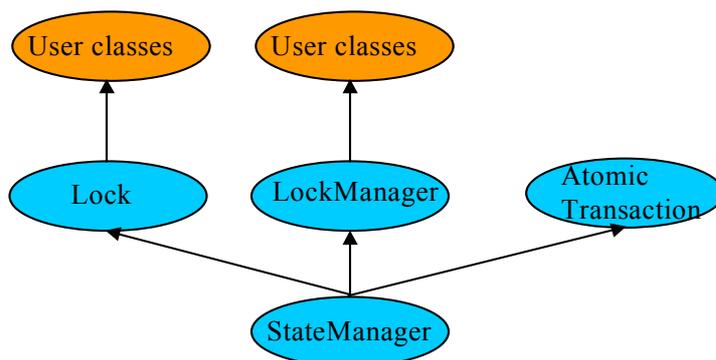


Figure 1: TxCore class hierarchy.

Apart from specifying the scopes of transactions, and setting appropriate locks within objects, the application programmer does not have any other responsibilities: *TxCore* and *Transactional Objects for Java (TXOJ)* guarantee that transactional objects will be registered with, and be driven by, the appropriate transactions, and crash recovery mechanisms are invoked automatically in the event of failures.

Saving object states

TxCore needs to be able to remember the state of an object for several purposes, including recovery (the state represents some past state of the object) and persistence (the state represents the final state of an object at application termination). Since these requirements have common functionality they are all implemented using the same mechanism: the classes `InputObjectState` and `OutputObjectState`. The classes maintain an internal array into which instances of the standard types can be contiguously packed (unpacked) using appropriate `pack` (`unpack`) operations. This buffer is automatically resized as required should it have insufficient space. The instances are all stored in the buffer in a standard form (so-called network byte order) to make them machine independent. Any other architecture independent format (such as XDR or ASN.1) could be implemented simply by replacing the operations with ones appropriate to the encoding required.

The object store

Implementations of persistence can be affected by restrictions imposed by the Java `SecurityManager`. Therefore, the object store provided with *TxCore* is implemented using the techniques of interface/implementation. The current distribution has implementations which write object states to the local file system or database, and remote implementations, where the interface uses a client stub (proxy) to remote services.

Persistent objects are assigned unique identifiers (instances of the `UId` class), when they are created, and this is used to identify them within the object store. States are read using the `read_committed` operation and written by the `write_(un)committed` operations.

Recovery and persistence

At the root of the class hierarchy is the class `StateManager`. This class is responsible for object activation and deactivation and object recovery. The simplified signature of the class is:

```
public abstract class StateManager
{
    public boolean activate ();
    public boolean deactivate (boolean commit);

    public Uid get_uid (); // object's identifier.

    // methods to be provided by a derived class

    public boolean restore_state (InputObjectState os);
    public boolean save_state (OutputObjectState os);

    protected StateManager ();
    protected StateManager (Uid id);
};
```

Objects are assumed to be of three possible flavours. They may simply be *recoverable*, in which case `StateManager` will attempt to generate and maintain appropriate recovery information for the object. Such objects have lifetimes that do not exceed the application program that creates them. Objects may be *recoverable and persistent*, in which case the lifetime of the object is assumed to be greater than that of the creating or accessing application, so that in addition to maintaining recovery information `StateManager` will attempt to automatically load (unload) any existing persistent state for the object by calling the `activate` (`deactivate`) operation at appropriate times. Finally, objects may possess none of these capabilities, in which case no recovery information is ever kept nor is object activation/deactivation ever automatically attempted.

If an object is *recoverable* or *recoverable and persistent* then `StateManager` will invoke the operations `save_state` (while performing `deactivate`), and `restore_state` (while performing `activate`) at various points during the execution of the application. These operations *must* be implemented by the programmer since `StateManager` cannot detect user level state changes. (We are examining the automatic generation of default `save_state` and `restore_state` operations, allowing the programmer to override this when application specific knowledge can be used to improve efficiency.) This gives the programmer the ability to decide which parts of an object's state should be made persistent. For example, for a spreadsheet it may not be necessary to save all entries if some values can simply be recomputed. The `save_state` implementation for a class `Example` that has integer member variables called A, B and C could simply be:

```
public boolean save_state(OutputObjectState o)
{
    if (!super.save_state(o))
        return false;

    try
    {
        o.packInt(A);
        o.packInt(B);
    }
```

```

        o.packInt(C);
    }
    catch (Exception e)
    {
        return false;
    }

    return true;
}

```

Note: it is necessary for all `save_state` and `restore_state` methods to call `super.save_state` and `super.restore_state`. This is to cater for improvements in the crash recovery mechanisms.

The life-cycle of a Transactional Object for Java

A persistent object not in use is assumed to be held in a *passive* state with its state residing in an object store and *activated* on demand. The fundamental life cycle of a persistent object in *TXOJ* is shown in Figure 2.

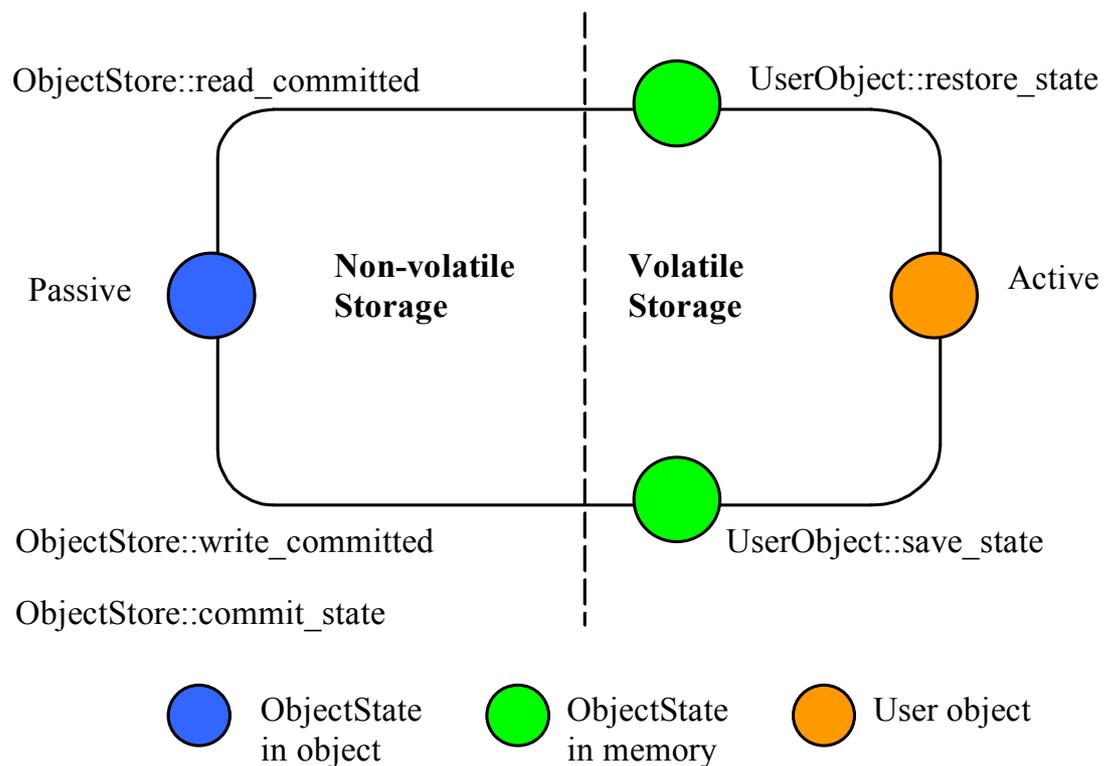


Figure 2: The life cycle of a persistent object.

- The object is initially passive, and is stored in the object store as an instance of the class `OutputObjectState`.

- When required by an application the object is automatically activated by reading it from the store using a `read_committed` operation and is then converted from an `InputObjectState` instance into a fully-fledged object by the `restore_state` operation of the object.
- When the application has finished with the object it is deactivated by converting it back into an `OutputObjectState` instance using the `save_state` operation, and is then stored back into the object store as a shadow copy using `write_uncommitted`. This shadow copy can be committed, overwriting the previous version, using the `commit_state` operation. The existence of shadow copies is normally hidden from the programmer by the transaction system. Object de-activation normally only occurs when the top-level transaction within which the object was activated commits.

Note: During its life time, a persistent object may be made active then passive many times.

The concurrency controller

The concurrency controller is implemented by the class `LockManager` which provides sensible default behaviour while allowing the programmer to override it if deemed necessary by the particular semantics of the class being programmed. As with `StateManager` and persistence, concurrency control implementations are accessed through interfaces. As well as providing access to remote services, the current implementations of concurrency control available to interfaces include:

- local disk/database implementation, where locks are made persistent by being written to the local file system or database.
- a purely local implementation, where locks are maintained within the memory of the virtual machine which created them; this implementation has better performance than when writing locks to the local disk, but objects cannot be shared between virtual machines. Importantly, it is a basic Java object with no requirements which can be affected by the `SecurityManager`.

The primary programmer interface to the concurrency controller is via the `setlock` operation. By default, the runtime system enforces strict two-phase locking following a multiple reader, single writer policy on a per object basis. However, as shown in Figure 1, by inheriting from the `Lock` class it is possible for programmers to provide their own lock implementations with different lock conflict rules to enable *type specific concurrency control*.

Lock acquisition is (of necessity) under programmer control, since just as `StateManager` cannot determine if an operation modifies an object, `LockManager` cannot determine if an operation requires a read or write lock. Lock release, however, is under control of the system and requires no further intervention by the programmer. This ensures that the two-phase property can be correctly maintained.

```
public abstract class LockManager extends StateManager
{
    public LockResult setlock (Lock toSet, int retry, int timeout);
};
```

The `LockManager` class is primarily responsible for managing requests to set a lock on an object or to release a lock as appropriate. However, since it is derived from `StateManager`, it can also control when some of the inherited facilities are invoked. For example, `LockManager` assumes that the setting of a write lock implies that the invoking operation must be about to modify the object. This may in turn cause recovery information to be saved if the object is recoverable. In a similar fashion, successful lock acquisition causes `activate` to be invoked.

The code below shows how we may try to obtain a write lock on an object:

```
public class Example extends LockManager
{
public boolean foobar ()
{
    AtomicAction A = new AtomicAction;
    boolean result = false;

    A.begin();

    if (setlock(new Lock(LockMode.WRITE), 0) == Lock.GRANTED)
    {
        /*
         * Do some work, and TXOJ will
         * guarantee ACID properties.
         */

        // automatically aborts if fails

        if (A.commit() == AtomicAction.COMMITTED)
        {
            result = true;
        }
    }
    else
        A.rollback();

    return result;
}
}
```

The transaction protocol engine

The transaction protocol engine is represented by the `AtomicAction` class, which uses `StateManager` in order to record sufficient information for crash recovery mechanisms to complete the transaction in the event of failures. It has methods for starting and terminating the transaction, and, for those situations where programmers require to implement their own resources, methods for registering them with the current transaction. Because *TxCore* supports subtransactions, if a transaction is begun within the scope of an already executing transaction it will automatically be nested.

Note: *TxCore* is multi-threaded aware, allowing each thread within an application to share a transaction or execute within its own transaction. Therefore, all *TxCore* classes are also thread safe.

Example

The simple example below illustrates the relationships between activation, termination and commitment:

```
{
. . .
O1 objct1 = new objct1(Name-A); /* (i) bind to "old" persistent object A */
O2 objct2 = new objct2();      /* create a "new" persistent object */
OTS.current().begin();        /* (ii) start of atomic action */

objct1.op(...);              /* (iii) object activation and invocations */
objct2.op(...);

. . .
OTS.current().commit(true); /* (iv) tx commits & objects deactivated */
}                               /* (v) */
```

The execution of the above code involves the following sequence of activities:

1. Creation of bindings to persistent objects; this could involve the creation of stub objects and a call to remote objects. In the above example we re-bind to an existing persistent object identified by Name-A, and a new persistent object. A naming system for remote objects maintains the mapping between object names and locations and is described in a later chapter.
2. Start of the atomic transaction.
3. Operation invocations: as a part of a given invocation the object implementation is responsible to ensure that it is locked in read or write mode (assuming no lock conflict), and initialised, if necessary, with the latest committed state from the object store. The first time a lock is acquired on an object within a transaction the object's state is acquired, if possible, from the object store.
4. Commit of the top-level action. This includes updating of the state of any modified objects in the object store.
5. Breaking of the previously created bindings.

The class hierarchy

The principal classes which make up the class hierarchy of *TxCore* are depicted below.

```
StateManager // Basic naming, persistence and recovery control
LockManager  // Basic two-phase locking concurrency control service
User-Defined Classes
Lock         // Standard lock type for multiple readers/single writer
User-Defined Lock Classes
AbstractRecord // Important utility class, similar to Resource
RecoveryRecord // handles object recovery
LockRecord    // handles object locking
RecordList    // Intentions list
other management record types
AtomicAction  // Implements transaction control abstraction
TopLevelTransaction
Input/OutputBuffer // Architecture neutral representation of an objects'
state
Input/OutputObjectState // Convenient interface to Buffer
ObjectStore     // Interface to the object storage services
```

Figure 3: TXOJ class hierarchy.

Programmers of fault-tolerant applications will be primarily concerned with the classes `LockManager`, `Lock` and `AtomicAction`. Other classes important to a programmer are `Uid`, and `ObjectState`. Most *TxCORE* classes are derived from the base class `StateManager`, which provides primitive facilities necessary for managing persistent and recoverable objects. These facilities include support for the activation and de-activation of objects, and state-based object recovery. The class `LockManager` uses the facilities of `StateManager` and `Lock` to provide the concurrency control (two-phase locking in the current implementation) required for implementing the serialisability property of atomic actions. The implementation of atomic action facilities is supported by `AtomicAction` and `TopLevelTransaction`.

Most *TxCORE* system classes are derived from the base class `StateManager`, which provides primitive facilities necessary for managing persistent and recoverable objects. These facilities include support for the activation and de-activation of objects, and state-based object recovery. The class `LockManager` uses the facilities of `StateManager` and provides the concurrency control required for implementing the serialisability property of atomic actions.

Consider a simple example. Assume that `Example` is a user-defined persistent class suitably derived from the `LockManager`. An application containing an atomic transaction `Trans` accesses an object (called `O`) of type `Example` by invoking the operation `op1` which involves state changes to `O`. The serialisability property requires that a write lock must be acquired on `O` before it is modified; thus the body of `op1` should contain a call to the `setlock` operation of the concurrency controller:

```
public boolean op1 (...)
{
    if (setlock (new Lock(LockMode.WRITE) == LockResult.GRANTED)
    {
        // actual state change operations follow
        ...
    }
}
```

Program 1: Simple Concurrency Control.

The operation `setlock`, provided by the `LockManager` class, performs the following functions in this case:

1. Check write lock compatibility with the currently held locks, and if allowed:
2. Call the `StateManager` operation `activate` that will load, if not done already, the latest persistent state of `O` from the object store. Then call the `StateManager` operation `modified` which has the effect of creating an instance of either `RecoveryRecord` or `PersistenceRecord` for `O` depending upon whether `O` was persistent or not (the `Lock` is a `WRITE` lock so the old state of the object must be retained prior to modification) and inserting it into the `RecordList` of `Trans`.
3. Create and insert a `LockRecord` instance in the `RecordList` of `Trans`.

Now suppose that action `Trans` is aborted sometime after the lock has been acquired. Then the `rollback` operation of `AtomicAction` will process the `RecordList` instance associated with `Trans` by invoking an appropriate `Abort` operation on the various records.

The implementation of this operation by the `LockRecord` class will release the `WRITE` lock while that of `RecoveryRecord/PersistenceRecord` will restore the prior state of `O`.

It is important to realise that all of the above work is automatically being performed by *TxCore* on behalf of the application programmer. The programmer need only start the transaction and set an appropriate lock; *TxCore and Transactional Objects for Java* take care of participant registration, persistence, concurrency control and recovery.

Chapter 2

Using TxCore

Introduction

In this section we shall describe *TxCore* and *Transactional Objects for Java* in more detail, and show how it can be used to construct transactional applications.

State management

Object states

TxCore needs to be able to remember the state of an object for several purposes, including recovery (the state represents some past state of the object), and for persistence (the state represents the final state of an object at application termination). Since all of these requirements require common functionality they are all implemented using the same mechanism - the classes `Input/OutputObjectState` and `Input/OutputBuffer`.

OutputBuffer

```
public class OutputBuffer
{
    public OutputBuffer ();

    public final synchronized boolean valid ();
    public synchronized byte[] buffer();
    public synchronized int length ();

    /* pack operations for standard Java types */

    public synchronized void packByte (byte b) throws IOException;
    public synchronized void packBytes (byte[] b) throws IOException;
    public synchronized void packBoolean (boolean b) throws IOException;
    public synchronized void packChar (char c) throws IOException;
    public synchronized void packShort (short s) throws IOException;
    public synchronized void packInt (int i) throws IOException;
    public synchronized void packLong (long l) throws IOException;
    public synchronized void packFloat (float f) throws IOException;
    public synchronized void packDouble (double d) throws IOException;
    public synchronized void packString (String s) throws IOException;
};
```

InputBuffer

```
public class InputBuffer
{
    public InputBuffer ();
```

```
public final synchronized boolean valid ();
public synchronized byte[] buffer();
public synchronized int length ();

    /* unpack operations for standard Java types */

public synchronized byte unpackByte () throws IOException;
public synchronized byte[] unpackBytes () throws IOException;
public synchronized boolean unpackBoolean () throws IOException;
public synchronized char unpackChar () throws IOException;
public synchronized short unpackShort () throws IOException;
public synchronized int unpackInt () throws IOException;
public synchronized long unpackLong () throws IOException;
public synchronized float unpackFloat () throws IOException;
public synchronized double unpackDouble () throws IOException;
public synchronized String unpackString () throws IOException;
};
```

The `Input/OutputBuffer` class maintains an internal array into which instances of the standard Java types can be contiguously packed (unpacked) using the `pack` (`unpack`) operations. This buffer is automatically resized as required should it have insufficient space. The instances are all stored in the buffer in a standard form (so-called network byte order) to make them machine independent.

OutputObjectState

```
class OutputObjectState extends OutputBuffer
{
public OutputObjectState (Uid newUid, String typeName);

public boolean notempty ();
public int size ();
public Uid stateUid ();
public String type ();
};
```

InputObjectState

```
class InputObjectState extends InputBuffer
{
public OutputObjectState (Uid newUid, String typeName, byte[] b);

public boolean notempty ();
public int size ();
public Uid stateUid ();
public String type ();
};
```

The class `Input/OutputObjectState` provides all the functionality of `Input/OutputBuffer` (through inheritance) but adds two additional instance variables that signify the `Uid` and `type` of the object for which the `Input/OutputObjectState` instance is a compressed image. These are used when accessing the object store during storage and retrieval of the object state.

The object store

The object store provided with *TxCore* deliberately has a fairly restricted interface so that it can be implemented in a variety of ways. For example, object stores are implemented in shared memory; on the Unix file system (in several different forms); and as a remotely accessible store. More complete information about the object stores available in *TxCore* can be found in the *Appendix*.

Note: as with all *TxCore* classes the default object stores are pure Java implementations; to access the shared memory and other more complex object store implementations it is necessary to use native methods.

All of the object stores hold and retrieve instances of the class `Input/OutputObjectState`. These instances are named by the `Uid` and `Type` of the object that they represent. States are read using the `read_committed` operation and written by the system using the `write_uncommitted` operation. Under normal operation new object states do not overwrite old object states but are written to the store as *shadow copies*. These shadows replace the original only when the `commit_state` operation is invoked. Normally all interaction with the object store is performed by *TxCore* system components as appropriate thus the existence of any shadow versions of objects in the store are hidden from the programmer.

```
public class ObjectStore
{
    public static final int OS_COMMITTED;
    public static final int OS_UNCOMMITTED;
    public static final int OS_COMMITTED_HIDDEN;
    public static final int OS_UNCOMMITTED_HIDDEN;
    public static final int OS_UNKNOWN;

    /* The abstract interface */
    public abstract boolean commit_state (Uid u, String name)
        throws ObjectStoreException;
    public abstract InputObjectState read_committed (Uid u, String name)
        throws ObjectStoreException;
    public abstract boolean write_uncommitted (Uid u, String name,
        OutputObjectState os) throws ObjectStoreException;
    . . .
};
```

When a transactional object is committing it is necessary for it to make certain state changes persistent in order that it can recover in the event of a failure and either continue to commit, or rollback. When using Transactional Objects for Java, *TxCore* will take care of this automatically. To guarantee ACID properties, these state changes must be *flushed* to the persistence store implementation before the transaction can proceed to commit; if they are not, the application may assume that the transaction has committed when in fact the state changes may still reside within an operating system cache, and may be lost by a subsequent machine failure. By default, *TxCore* ensures that such state changes are flushed. However, doing so can impose a significant performance penalty on the application. To prevent transactional object state flushes, set the `com.arjuna.ats.arjuna.objectstore.objectStoreSync` variable to `OFF`.

Selecting an object store implementation

TxCore comes with support for several different object store implementations. The *Appendix* describes these implementations, how to select and configure a given implementation (using the `com.arjuna.ats.arjuna.objectstore.objectStoreType` property variable) on a per object basis, and indicates how additional implementations can be provided.

StateManager

The *TxCore* class `StateManager` manages the state of an object and provides all of the basic support mechanisms required by an object for state management purposes. `StateManager` is responsible for creating and registering appropriate resources concerned with the persistence and recovery of the transactional object. If a transaction is nested, then `StateManager` will also propagate these resources between child transactions and their parents at commit time.

Objects in *TxCore* are assumed to be of three possible basic flavours. They may simply be *recoverable*, in which case `StateManager` will attempt to generate and maintain appropriate recovery information for the object (as instances of the class `Input/OutputObjectState`). Such objects have lifetimes that do not exceed the application program that creates them. Objects may be *recoverable and persistent*, in which case the lifetime of the object is assumed to be greater than that of the creating or accessing application so that in addition to maintaining recovery information `StateManager` will attempt to automatically load (unload) any existing persistent state for the object by calling the `activate` (`deactivate`) operation at appropriate times. Finally, objects may possess none of these capabilities in which case no recovery information is ever kept nor is object activation/deactivation ever automatically attempted. This object property is selected at object construction time and cannot be changed thereafter. Thus an object cannot gain (or lose) recovery capabilities at some arbitrary point during its lifetime.

```
public class ObjectStatus
{
    public static final int PASSIVE;
    public static final int PASSIVE_NEW;
    public static final int ACTIVE;
    public static final int ACTIVE_NEW;
    public static final int UNKNOWN_STATUS;
};

public class ObjectType
{
    public static final int RECOVERABLE;
    public static final int ANDPERSISTENT;
    public static final int NEITHER;
};

public abstract class StateManager
{
    public synchronized boolean activate ();
    public synchronized boolean activate (String storeRoot);
    public synchronized boolean deactivate ();
    public synchronized boolean deactivate (String storeRoot, boolean commit);

    public synchronized void destroy ();
};
```

```

public final Uid get_uid ();

public boolean restore_state (InputObjectState, int ObjectType);
public boolean save_state (OutputObjectState, int ObjectType);
public String type ();
    . . .

protected StateManager ();
protected StateManager (int ObjectType, ObjectName attr);
protected StateManager (Uid uid);
protected StateManager (Uid uid, ObjectName attr);
    . . .

protected final void modified ();
    . . .
};

public class ObjectModel
{
public static final int SINGLE;
public static final int MULTIPLE;
};

```

If an object is recoverable (or persistent) then `StateManager` will invoke the operations `save_state` (while performing deactivation), `restore_state` (while performing activate) and `type` at various points during the execution of the application. These operations *must* be implemented by the programmer since `StateManager` does not have access to a runtime description of the layout of an arbitrary Java object in memory and thus cannot implement a default policy for converting the in memory version of the object to its passive form. However, the capabilities provided by `Input/OutputObjectState` make the writing of these routines fairly simple. For example, the `save_state` implementation for a class `Example` that had member variables called `A`, `B` and `C` could simply be the following:

```

public boolean save_state ( OutputObjectState os, int ObjectType )
{
    if (!super.save_state(os, ObjectType))
        return false;

    try
    {
        os.packInt(A);
        os.packString(B);
        os.packFloat(C);

        return true;
    }
    catch (IOException e)
    {
        return false;
    }
}

```

In order to support crash recovery for persistent objects it is necessary for all `save_state` and `restore_state` methods of user objects to call `super.save_state` and `super.restore_state`.

Note: The `type` method is used to determine the location in the object store where the state of instances of that class will be saved and ultimately restored. This can actually be any valid string. However, you should avoid using the hash character (`#`) as this is reserved for special directories that TxCore requires.

The `get_uid` operation of `StateManager` provides read only access to an object's internal system name for whatever purpose the programmer requires (such as registration of the name in a name server). The value of the internal system name can *only* be set when an object is initially constructed - either by the provision of an explicit parameter or by generating a new identifier when the object is created.

The `destroy` method can be used to remove the object's state from the object store. This is an atomic operation, and therefore will only remove the state if the top-level transaction within which it is invoked eventually commits. The programmer *must* obtain exclusive access to the object prior to invoking this operation.

Since object recovery and persistence essentially have complimentary requirements (the only difference being where state information is stored and for what purpose) `StateManager` effectively combines the management of these two properties into a single mechanism. That is, it uses instances of the class `Input/OutputObjectState` both for recovery and persistence purposes. An additional argument passed to the `save_state` and `restore_state` operations allows the programmer to determine the purpose for which any given invocation is being made thus allowing different information to be saved for recovery and persistence purposes.

Object models

TxCore supports two models for objects, which as we shall show affect how an objects state and concurrency control are implemented:

- **SINGLE:** only a single copy of the object exists within the application; this will reside within a single JVM, and all clients must address their invocations to this server. This model provides better performance, but represents a single point of failure, and in a multi-threaded environment may not protect the object from corruption if a single thread fails.

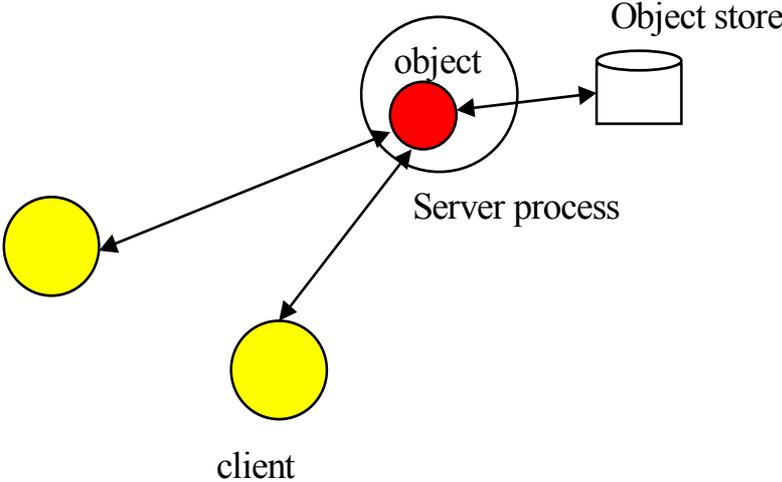


Figure 4 SINGLE object model.

- MULTIPLE: logically a single instance of the object exists, but copies of it are distributed across different JVMs; the performance of this model is worse than the SINGLE model, but it provides better failure isolation.

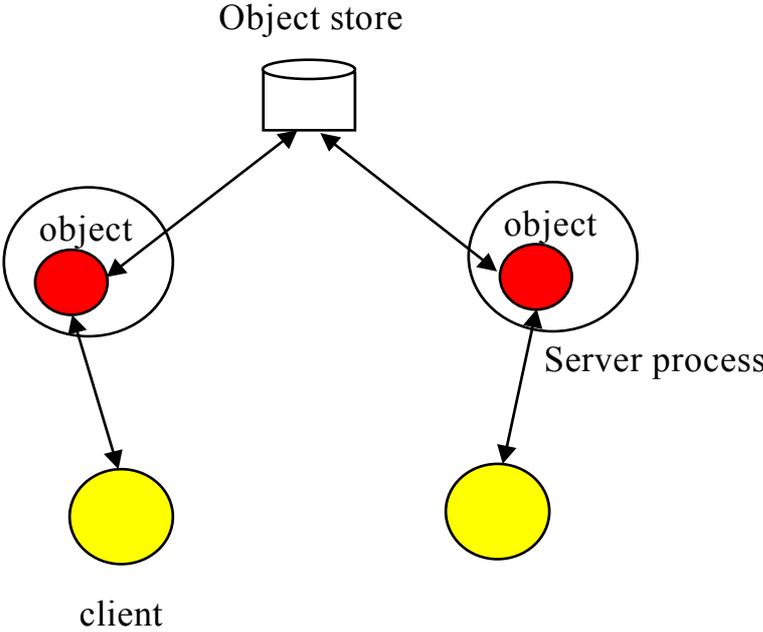


Figure 5 MULTIPLE object model.

The default model is SINGLE. The programmer can override this on a per object basis by providing an appropriate instance of the `com.arjuna.ats.arjuna.gandiva.ObjectName` class at object construction.

Note: The model can be changed between each successive instantiation of the object, i.e., it need not be the same during the object's lifetime.

To provide a suitable `ObjectName` class, it is necessary to perform the following steps:

- create a new instance of `ObjectName`.
- set the object model attribute using the `com.arjuna.ats.arjuna.ArjunaNames.StateManager_objectModel()` name.

For example:

```
{
    ObjectName attr = new ObjectName("SNS:myObjectName");

    attr.setLongAttribute(ArjunaNames.StateManager_objectModel(),
                        ObjectModel.SINGLE);

    AtomicObject obj = new AtomicObject(ObjectType.ANDPERSISTENT, attr);
}
```

Summary

In summary, the *TxCore* class `StateManager` manages the state of an object and provides all of the basic support mechanisms required by an object for state management purposes. Some operations *must* be defined by the class developer. These operations are: `save_state`, `restore_state`, and `type`.

boolean save_state (OutputObjectState state, int ObjectType)

Invoked whenever the state of an object might need to be saved for future use - primarily for recovery or persistence purposes. The `ObjectType` parameter indicates the reason that `save_state` was invoked by *TxCore*. This enables the programmer to save different pieces of information into the `OutputObjectState` supplied as the first parameter depending upon whether the state is needed for recovery or persistence purposes. For example, pointers to other *TxCore* objects might be saved simply as pointers for recovery purposes but as `Uid`'s for persistence purposes. As shown earlier, the `OutputObjectState` class provides convenient operations to allow the saving of instances of all of the basic types in Java. In order to support crash recovery for persistent objects it is necessary for all `save_state` methods to call `super.save_state`.

Note: `save_state` assumes that an object is internally consistent and that all variables saved have valid values. It is the programmer's responsibility to ensure that this is the case.

boolean restore_state (InputObjectState state, int ObjectType)

Invoked whenever the state of an object needs to be restored to the one supplied. Once again the second parameter allows different interpretations of the supplied state. In order to support crash recovery for persistent objects it is necessary for all `restore_state` methods to call `super.restore_state`.

String type ()

The *TxCore* persistence mechanism *requires* a means of determining the type of an object as a string so that it can save/restore the state of the object into/from the object store. By convention this information indicates the position of the class in the hierarchy. For example, `"/StateManager/LockManager/Object"`.

Caution: The type method is used to determine the location in the object store where the state of instances of that class will be saved and ultimately restored. This can actually be any valid string. However, you should avoid using the hash character (#) as this is reserved for special directories that TxCore requires.

Example

Consider the following basic `Array` class derived from the `StateManager` class (in this example, to illustrate saving and restoring of an object's state, the `highestIndex` variable is used to keep track of the highest element of the array that has a non-zero value):

```
public class Array extends StateManager
{
    public Array ();
    public Array (Uid objUid);
    public void finalize ( super.terminate(); );

    /* Class specific operations. */

    public boolean set (int index, int value);
    public int get (int index);

    /* State management specific operations. */

    public boolean save_state (OutputObjectState os, int ObjectType);
    public boolean restore_state (InputObjectState os, int ObjectType);
    public String type ();

    public static final int ARRAY_SIZE = 10;

    private int[] elements = new int[ARRAY_SIZE];
    private int highestIndex;
};
```

The `save_state`, `restore_state` and `type` operations can be defined as follows:

```
/* Ignore ObjectType parameter for simplicity */

public boolean save_state (OutputObjectState os, int ObjectType)
{
    if (!super.save_state(os, ObjectType))
        return false;

    try
    {
        packInt (highestIndex);
    }

    /*
```

```
    * Traverse array state that we wish to save. Only save active elements
    */

    for (int i = 0; i <= highestIndex; i++)
        os.packInt(elements[i]);

        return true;
    }
    catch (IOException e)
    {
        return false;
    }
}
public boolean restore_state (InputObjectState os, int ObjectType)
{
    if (!super.restore_state(os, ObjectType))
        return false;

    try
    {
        int i = 0;

        highestIndex = os.unpackInt();

        while (i < ARRAY_SIZE)
        {
            if (i <= highestIndex)
                elements[i] = os.unpackInt();
            else
                elements[i] = 0;
            i++;
        }

        return true;
    }
    catch (IOException e)
    {
        return false;
    }
}
public String type ()
{
    return "/StateManager/Array";
}
```

Lock management and concurrency control

Concurrency control information within *TxCORE* is maintained by *locks*. Locks which are required to be shared between objects in different processes may be held within a lock store, similar to the object store facility presented previously. The lock store provided with *TxCORE* deliberately has a fairly restricted interface so that it can be implemented in a variety of ways. For example, lock stores are implemented in shared memory; on the Unix file system (in several different forms); and as a remotely accessible store. More information about the object stores available in *TxCORE* can be found in the *Appendix*.

Note: as with all *TxCORE* classes the default lock stores are pure Java implementations; to access the shared memory and other more complex lock store implementations it is necessary to use native methods.

```

public class LockStore
{
    public abstract InputObjectState read_state (UId u, String tName)
                                           throws
                                           LockStoreException;

    public abstract boolean remove_state (UId u, String tname);
    public abstract boolean write_committed (UId u, String tName,
                                           OutputObjectState state);
};

```

Selecting a lock store implementation

TxCore comes with support for several different object store implementations. If the object model being used is SINGLE, then no lock store is required for maintaining locks, since the information about the object is not exported from it. However, if the MULTIPLE model is used, then different run-time environments (processes, Java virtual machines) may need to share concurrency control information. The implementation type of the lock store to use can be specified for all objects within a given execution environment using the `com.arjuna.ats.txoj.lockstore.lockStoreType` property variable. Currently this can have one of the following values:

- *BasicLockStore*: this is an in-memory implementation which does not, by default, allow sharing of stored information between execution environments. The application programmer is responsible for sharing the store information.
- *BasicPersistentLockStore*: this is the default implementation, and stores locking information within the local file system. Therefore execution environments *that share the same file store* can share concurrency control information. The root of the file system into which locking information is written is the `LockStore` directory within the *TxCore* installation directory. This can be overridden at runtime by setting the `com.arjuna.ats.txoj.lockstore.lockStoreDir` property variable accordingly, or placing the location within the CLASSPATH:

```

java -D com.arjuna.ats.txoj.lockstore.lockStoreDir=/var/tmp/LockStore
myprogram

```

or

```

java -classpath $CLASSPATH;/var/tmp/LockStore myprogram

```

If neither of these approaches is taken, then the default location will be at the same level as the `etc` directory of the installation.

LockManager

The concurrency controller is implemented by the class `LockManager` which provides sensible default behaviour while allowing the programmer to override it if deemed necessary by the particular semantics of the class being programmed. The primary programmer interface to the concurrency controller is via the `setlock` operation. By default, the *TxCore* runtime system enforces strict two-phase locking following a multiple reader, single writer policy on a per object basis. Lock acquisition is under programmer control, since just as

`StateManager` cannot determine if an operation modifies an object, `LockManager` cannot determine if an operation requires a read or write lock. Lock release, however, is normally under control of the system and requires no further intervention by the programmer. This ensures that the two-phase property can be correctly maintained.

The `LockManager` class is primarily responsible for managing requests to set a lock on an object or to release a lock as appropriate. However, since it is derived from `StateManager`, it can also control when some of the inherited facilities are invoked. For example, if a request to set a write lock is granted, then `LockManager` invokes `modified` directly assuming that the setting of a write lock implies that the invoking operation must be about to modify the object. This may in turn cause recovery information to be saved if the object is recoverable. In a similar fashion, successful lock acquisition causes `activate` to be invoked.

Therefore, `LockManager` is directly responsible for activating/de-activating persistent objects, and registering `Resources` for managing concurrency control. By driving the `StateManager` class, it is also responsible for registering `Resources` for persistent/recoverable state manipulation and object recovery. The application programmer simply sets appropriate locks, starts and ends transactions, and extends the `save_state` and `restore_state` methods of `StateManager`.

```
public class LockResult
{
    public static final int GRANTED;
    public static final int REFUSED;
    public static final int RELEASED;
};

public class ConflictType
{
    public static final int CONFLICT;
    public static final int COMPATIBLE;
    public static final int PRESENT;
};

public abstract class LockManager extends StateManager
{
    public static final int defaultTimeout;
    public static final int defaultRetry;
    public static final int waitTotalTimeout;

    public synchronized int setlock (Lock l);
    public synchronized int setlock (Lock l, int retry);
    public synchronized int setlock (Lock l, int retry, int sleepTime);
    public synchronized boolean releaselock (Uid uid);

    /* abstract methods inherited from StateManager */

    public boolean restore_state (InputObjectState os, int ObjectType);
    public boolean save_state (OutputObjectState os, int ObjectType);
    public String type ();

    protected LockManager ();
    protected LockManager (int ObjectType, ObjectName attr);
    protected LockManager (Uid storeUid);
    protected LockManager (Uid storeUid, int ObjectType, ObjectName attr);
    . . .
};
```

The `setlock` operation must be parameterised with the type of lock required (`READ / WRITE`), and the number of retries to acquire the lock before giving up. If a lock conflict occurs, one of the following scenarios will take place:

If the retry value is equal to `LockManager.waitTotalTimeout`, then the thread which called `setlock` will be blocked until the lock is released, or the total timeout specified has elapsed, and in which `REFUSED` will be returned.

If the lock cannot be obtained initially then `LockManager` will try for the specified number of retries, waiting for the specified timeout value between each failed attempt. The default is 100 attempts, each attempt being separated by a 0.25 seconds delay; the time between retries is specified in *micro-seconds*.

If a lock conflict occurs the current implementation simply times out lock requests, thereby preventing deadlocks, rather than providing a full deadlock detection scheme. If the requested lock is obtained, the `setlock` operation will return the value `GRANTED`, otherwise the value `REFUSED` is returned. It is the responsibility of the programmer to ensure that the remainder of the code for an operation is only executed if a lock request is granted. Below are examples of the use of the `setlock` operation.

```
res = setlock(new Lock(WRITE), 10); // Will attempt to set a
    // write lock 11 times (10
    // retries) on the object
    // before giving up.
res = setlock(new Lock(READ), 0); // Will attempt to set a read
    // lock 1 time (no retries) on
    // the object before giving up.
res = setlock(new Lock(WRITE); // Will attempt to set a write
    // lock 101 times (default of
    // 100 retries) on the object
    // before giving up.
```

The concurrency control mechanism is integrated into the atomic action mechanism, thus ensuring that as locks are granted on an object appropriate information is registered with the currently running atomic action to ensure that the locks are released at the correct time. This frees the programmer from the burden of explicitly freeing any acquired locks if they were acquired within atomic actions. However, if locks are acquired on an object *outside* of the scope of an atomic action, it is the programmer's responsibility to release the locks when required, using the corresponding `releaselock` operation.

Locking policy

Unlike many other systems, locks in *TxCore* are not special system types. Instead they are simply instances of other *TxCore* objects (the class `Lock` which is also derived from `StateManager` so that locks may be made persistent if required and can also be named in a simple fashion). Furthermore, `LockManager` deliberately has no knowledge of the semantics of the actual policy by which lock requests are granted. Such information is maintained by the actual `Lock` class instances which provide operations (the `conflictsWith` operation) by which `LockManager` can determine if two locks conflict or not. This separation is important in that it allows the programmer to derive new lock types from the basic `Lock` class and by

providing appropriate definitions of the conflict operations enhanced levels of concurrency may be possible.

```
public class LockMode
{
    public static final int READ;
    public static final int WRITE;
};

public class LockStatus
{
    public static final int LOCKFREE;
    public static final int LOCKHELD;
    public static final int LOCKRETAINED;
};

public class Lock extends StateManager
{
    public Lock (int lockMode);

    public boolean conflictsWith (Lock otherLock);
    public boolean modifiesObject ();

    public boolean restore_state (InputObjectState os, int ObjectType);
    public boolean save_state (OutputObjectState os, int ObjectType);
    public String type ();
    . . .
};
```

The `Lock` class provides a `modifiesObject` operation which `LockManager` uses to determine if granting this locking request requires a call on `modified`. This operation is provided so that locking modes other than simple read and write can be supported. The supplied `Lock` class supports the traditional multiple reader/single writer policy.

Object construction and destruction

Recall that *TxCore* objects can be recoverable; recoverable and persistent; or neither. Additionally each object possesses a unique internal name. These attributes can *only* be set when that object is constructed. Thus `LockManager` provides two *protected* constructors for use by derived classes, each of which fulfils a distinct purpose:

LockManager ():

This constructor allows the creation of *new* objects, that is, no prior state is assumed to exist.

LockManager (int ObjectType, ObjectName attr):

As above, this constructor allows the creation of *new* objects, that is, no prior state is assumed to exist. The `ObjectType` parameter determines whether an object is simply recoverable (indicated by `RECOVERABLE`); recoverable and persistent (indicated by `ANDPERSISTENT`) or neither (`NEITHER`). If an object is marked as being persistent then the state of the object will be stored in one of the object stores. The `shared` parameter only has meaning if `ot` is `RECOVERABLE`; if `attr` is not null and the object model is `SINGLE` (the default behaviour) then the recoverable state of the object is maintained within the object itself (i.e., it has no

external representation), otherwise an in-memory (volatile) object store is used to store the state of the object between atomic actions.

Constructors for new persistent objects should make use of atomic actions within themselves. This will ensure that the state of the object is automatically written to the object store either when the action in the constructor commits or, if an enclosing action exists, when the appropriate top-level action commits. Later examples in this chapter illustrate this point further.

LockManager(Uid objUid):

This constructor allows access to an *existing* persistent object, whose internal name is given by the `objUid` parameter. Objects constructed using this operation will normally have their prior state (identified by `objUid`) loaded from an object store automatically by the system.

LockManager(Uid objUid, ObjectName attr):

As above, this constructor allows access to an *existing* persistent object, whose internal name is given by the `objUid` parameter. Objects constructed using this operation will normally have their prior state (identified by `objUid`) loaded from an object store automatically by the system. If the `attr` parameter is not null, and the object model is `SINGLE` (the default behaviour), then the object will not be reactivated at the start of each top-level transaction.

The destructor of a programmer-defined class *must* invoke the inherited operation `terminate` to inform the state management mechanism that the object is about to be destroyed otherwise unpredictable results may occur.

Because `LockManager` inherits from `StateManager`, it will pass any supplied `ObjectName` instance to the `StateManager` class. As such, it is possible to set the `StateManager` object model as described earlier.

General transaction issues

Advanced transaction issues with TxCore

Atomic actions (transactions) can be used by both application programmers and class developers. Thus entire operations (or parts of operations) can be made atomic as required by the semantics of a particular operation. This chapter will describe some of the more subtle issues involved with using transactions in general and *TxCore* in particular.

Checking transactions

In a multi-threaded application, multiple threads may be associated with a transaction during its lifetime, i.e., the thread's share the context. In addition, it is possible that if one thread terminates a transaction other threads may still be active within it. In a distributed environment, it can be difficult to guarantee that all threads have finished with a transaction when it is terminated. By default, *TxCore* will issue a warning if a thread terminates a transaction when other threads are still active within it; however, it will allow the transaction termination to continue. Other solutions to this problem are possible, e.g., blocking the thread which is terminating the transaction until all other threads have disassociated themselves from the transaction context. Therefore, *TxCore* provides the `com.arjuna.ats.arjuna.coordinator.CheckedAction` class, which allows the thread/transaction termination policy to be overridden. Each transaction has an instance of this class associated with it, and application programmers can provide their own implementations on a per transaction basis.

```
public class CheckedAction
{
    public CheckedAction ();

    public synchronized void check (boolean isCommit, Uid actUid,
                                     BasicList list);
};
```

When a thread attempts to terminate the transaction and there are active threads within it, the system will invoke the `check` method on the transaction's `CheckedAction` object. The parameters to the `check` method are:

- `isCommit`: indicates whether the transaction is in the process of committing or rolling back.
- `actUid`: the transaction identifier.

- `list`: a list of all of the threads currently marked as active within this transaction.

When `check` returns, the transaction termination will continue. Obviously the state of the transaction at this point may be different from that when `check` was called, e.g., the transaction may subsequently have been committed.

Statistics gathering

By default, the *JBossTS* does not maintain any history information about transactions. However, by setting the `com.arjuna.ats.arjuna.coordinator.enableStatistics` property variable to `YES`, the transaction service will maintain information about the number of transactions created, and their outcomes. This information can be obtained during the execution of a transactional application via the `com.arjuna.TxCore.Atomic.TxStats` class:

```
public class TxStats
{
    /**
     * Returns the number of transactions (top-level and nested)
     * created so far.
     */
    public static int numberOfTransactions ();

    /**
     * Returns the number of nested (sub) transactions created so far.
     */
    public static int numberOfNestedTransactions ();

    /**
     * Returns the number of transactions which have terminated with
     * heuristic outcomes.
     */
    public static int numberOfHeuristics ();

    /**
     * Returns the number of committed transactions.
     */
    public static int numberOfCommittedTransactions ();

    /**
     * Returns the number of transactions which have rolled back.
     */
    public static int numberOfAbortedTransactions ();
}
```

Last resource commit optimisation

In some cases it may be necessary to enlist participants that aren't two-phase commit aware into a two-phase commit transaction. If there is only a single resource then there is no need

for two-phase commit. However, what if there are multiple resources in the transaction? In this case, the Last Resource Commit optimization (LRCO) comes into play. It is possible for a single resource that is one-phase aware (i.e., can only commit or roll back, with no prepare), to be enlisted in a transaction with two-phase commit aware resources. The coordinator treats the one-phase aware resource slightly differently, in that it executes the prepare phase on all other resource first, and if it then intends to commit the transaction it passes control to the one-phase aware resource. If it commits, then the coordinator logs the decision to commit and attempts to commit the other resources as well.

In order to utilise the LRCO, your participant must implement the `com.arjuna.ats.arjuna.coordinator.OnePhase` interface and be registered with the transaction through the `BasicAction.add` operation; since this operation expects instances of `AbstractRecord`, you must create an instance `com.arjuna.ats.arjuna.LastResourceRecord` and give your participant as the constructor parameter, as shown below:

```
try
{
    boolean success = false;
    AtomicAction A = new AtomicAction();
    OnePhase opRes = new OnePhase(); // used OnePhase interface

    System.err.println("Starting top-level action.");

    A.begin();
    A.add(new LastResourceRecord(opRes));
    A.add(new ShutdownRecord(ShutdownRecord.FAIL_IN_PREPARE));

    A.commit();
}
```

Nested transactions

There are no special constructs for nesting of transactions: if an action is begun while another action is running then it is automatically nested. This allows for a modular structure to applications, whereby objects can be implemented using atomic actions within their operations without the application programmer having to worry about the applications which use them, i.e., whether or not the applications will use atomic actions as well. Thus, in some applications actions may be top-level, whereas in others they may be nested. Objects written in this way can then be shared between application programmers, and *TxCore* will guarantee their consistency.

If a nested action is aborted then all of its work will be undone, although strict two-phase locking means that any locks it may have obtained will be retained until the top-level action commits or aborts. If a nested action commits then the work it has performed will only be committed by the system if the top-level action commits; if the top-level action aborts then *all* of the work will be undone.

The committing or aborting of a nested action does not automatically affect the outcome of the action within which it is nested. This is application dependant, and allows a programmer to structure atomic actions to contain faults, undo work, etc.

Asynchronously committing a transaction

By default, *JBossTS* executes the commit protocol of a top-level transaction in a synchronous manner, i.e., all registered resources will be told to prepare in order by a single thread, and then they will be told to commit or rollback. This has several possible disadvantages:

- In the case of many registered resources, the prepare operation can logically be invoked in parallel on each resource. The disadvantage is that if an “early” resource in the list of registered resource forces a rollback during prepare, possibly many prepare operations will have been made needlessly.
- In the case where heuristic reporting is not required by the application, the second phase of the commit protocol can be done asynchronously, since its success or failure is not important.

Therefore, *JBossTS* provides runtime options to enable possible threading optimizations. By setting the `com.arjuna.ats.arjuna.coordinator.asyncPrepare` environment variable to `YES`, during the prepare phase a separate thread will be created for each registered participant within the transaction. By setting `com.arjuna.ats.arjuna.coordinator.asyncCommit` to `YES`, a separate thread will be created to complete the second phase of the transaction *if knowledge about heuristics outcomes is not required*.

Independent top-level transactions

In addition to normal top-level and nested atomic actions *TxCore* also supports *independent top-level* actions, which can be used to relax strict serialisability in a controlled manner. An independent top-level action can be executed from anywhere within another atomic action and behaves *exactly* like a normal top-level action, that is, its results are made permanent when it commits and will not be undone if any of the actions within which it was originally nested abort.

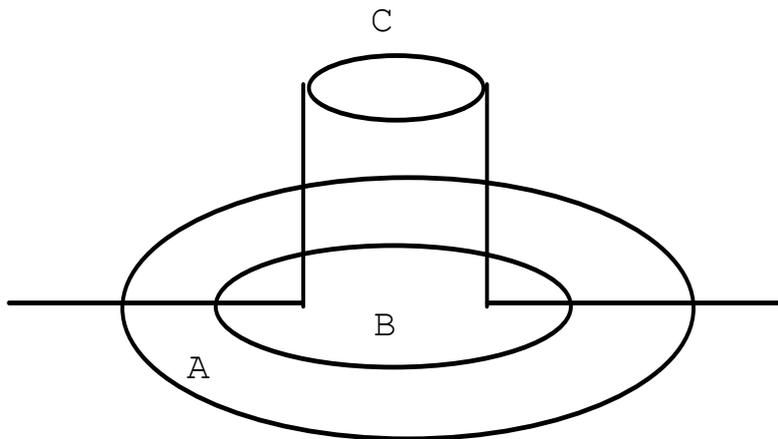


Figure 6 Independent Top-Level Action.

Figure 6 shows a typical nesting of atomic actions, where action B is nested within action A. Although atomic action C is logically nested within action B (it had its `Begin` operation invoked while B was active) because it is an independent top-level action, it will commit or abort independently of the other actions within the structure. Because of the nature of independent top-level actions they should be used with caution and only in situations where their use has been carefully examined.

Top-level actions can be used within an application by declaring and using instances of the class `TopLevelTransaction`. They are used in exactly the same way as other transactions.

Transactions within `save_state` and `restore_state`

Caution must be exercised when writing the `save_state` and `restore_state` operations to ensure that no atomic actions are started (either explicitly in the operation or implicitly through use of some other operation). This restriction arises due to the fact that `TxCore` may invoke `restore_state` as part of its commit processing resulting in the attempt to execute an atomic action during the commit or abort phase of another action. This might violate the atomicity properties of the action being committed (aborted) and is thus discouraged.

Example

If we consider the `Array` example given previously, the `set` and `get` operations could be implemented as shown below.

Note: this is a simplification of the code, ignoring error conditions and exceptions.

```
public boolean set (int index, int value)
{
    boolean result = false;
    AtomicAction A = new AtomicAction();

    A.begin();
```

```

// We need to set a WRITE lock as we want to modify the state.
if (setlock(new Lock(LockMode.WRITE), 0) == LockResult.GRANTED)
{
    elements[index] = value;
    if ((value > 0) &&(index > highestIndex)
        highestIndex = index;
    A.commit(true);
    result = true;
}
else
    A.rollback();

return result;
}

public int get (int index) // assume -1 means error
{
    AtomicAction A = new AtomicAction();

    A.begin();

    // We only need a READ lock as the state is unchanged.
    if (setlock(new Lock(LockMode.READ), 0) == LockResult.GRANTED)
    {
        A.commit(true);

        return elements[index];
    }
    else
        A.rollback();

    return -1;
}

```

Garbage collecting objects

Java objects are deleted when the garbage collector determines that they are no longer required. Deleting an object that is currently under the control of a transaction *must* be approached with caution since if the object is being manipulated within a transaction its fate is effectively determined by the transaction. Therefore, regardless of the references to a transactional object maintained by an application, *TxCore* will always retain its own references to ensure that the object is not garbage collected until after any transaction has terminated.

Transaction timeouts

By default transactions live until they are terminated by the application that created them or a failure occurs. However, it is possible to set a timeout (in seconds) on a per transaction basis such that if the transaction has not terminated before the timeout expires it will be automatically rolled back.

In *TxCore*, the timeout value is provided as a parameter to the `AtomicAction` constructor. If a value of `AtomicAction.NO_TIMEOUT` is provided (the default) then the transaction will

not be automatically timed out. Any other positive value is assumed to be the timeout for the transaction (in seconds). A value of zero is taken to be a global default timeout, which can be provided by the property `com.arjuna.ats.arjuna.coordinator.defaultTimeout`. Unless changed the default value is 60 seconds.

When a top-level transaction is created with a non-zero timeout, it is subject to being rolled back if it has not completed within the specified number of seconds. *JBossTS* uses a separate *reaper thread* which monitors all locally created transactions, and forces them to roll back if their timeouts elapse. To prevent this thread from consuming application time, it only runs periodically. The default checking period is 120000 *milliseconds*, but can be overridden by setting the `com.arjuna.ats.arjuna.coordinator.txReaperTimeout` property variable to another valid value, in microseconds. Alternatively, if the `com.arjuna.ats.arjuna.coordinator.txReaperMode` is set to `DYNAMIC`, the transaction reaper will wake whenever a transaction times out. This has the advantage of terminating transactions early, but may suffer from continually rescheduling the reaper thread.

If a value of 0 is specified for the timeout of a top-level transaction (or no timeout is specified), then *JBossTS* will not impose any timeout on the transaction, i.e., it will be allowed to run indefinitely. This default timeout can be overridden by setting the `com.arjuna.ats.arjuna.coordinator.defaultTimeout` property variable when using `ArjunaCore` or `ArjunaJTS`, or `com.arjuna.ats.jts.defaultTimeout` if using `ArjunaJTS`, to the required timeout value in *seconds*.

Chapter 4

Hints and tips

General

Using transactions in constructors

Examples throughout this manual have used transactions in the implementation of constructors for *new* persistent objects. This is deliberate because it guarantees correct propagation of the state of the object to the object store. Recall that the state of a modified persistent object is only written to the object store when the top-level transaction commits. Thus, if the constructor transaction is top-level and it commits, then the newly created object is written to the store and becomes available immediately. If however, the constructor transaction commits but is nested because some other transaction started prior to object creation is running, then the state will be written only if all of the parent transactions commit.

On the other hand, if the constructor does not use transactions then it is possible for inconsistencies in the system to arise. For example, if no transaction is active when the object is created then its state will not be saved to the store until the next time the object is modified under the control of some transaction.

Consider this simple example:

```
AtomicAction A = new AtomicAction();
Object obj1;
Object obj2;

obj1 = new Object();           // create new object
obj2 = new Object("old");     // existing object

A.begin(0);
obj2.remember(obj1.get_uid()); // obj2 now contains reference to obj1
A.commit(true);               // obj2 saved but obj1 is not
```

Here the two objects are created outside of the control of the top-level action A. `obj1` is a new object; `obj2` an old existing object. When the `remember` operation of `obj2` is invoked the object will be activated and the `Uid` of `obj1` remembered. Since this action commits the persistent state of `obj2` could now contain the `Uid` of `obj1`. However, the state of `obj1` itself has not been saved since it has not been manipulated under the control of any action. In fact, unless it is modified under the control of some action later in the application it will *never* be saved. If, however, the constructor had used an atomic action the state of `obj1` would have automatically been saved at the time it was constructed and this inconsistency could not arise.

More on `save_state` and `restore_state`

TxCore may invoke the user-defined `save_state` operation of an object effectively at any time during the lifetime of an object including during the execution of the *body* of the object's constructor (particularly if it uses atomic actions). It is important, therefore, that all of the variables saved by `save_state` are correctly initialised.

Caution must be also exercised when writing the `save_state` and `restore_state` operations to ensure that no transactions are started (either explicitly in the operation or implicitly through use of some other operation). This restriction arises due to the fact that *TxCore* may invoke `restore_state` as part of its commit processing resulting in the attempt to execute an atomic transaction during the commit or abort phase of another transaction. This might violate the atomicity properties of the transaction being committed (aborted) and is thus discouraged.

In order to support crash recovery for persistent objects it is necessary for all `save_state` and `restore_state` methods of user objects to call `super.save_state` and `super.restore_state`.

Packing objects

All of the basic types of Java (int, long, etc.) can be saved and restored from an `Input/OutputObjectState` instance by using the `pack` (and `unpack`) routines provided by `Input/OutputObjectState`. However packing and unpacking objects should be handled differently. This is because packing objects brings in the additional problems of *aliasing*. That is two different object references may in actual fact point at the same item. For example:

```
public class Test
{
    public Test (String s);
    ...
    private String s1;
    private String s2;
};

public Test (String s)
{
    s1 = s;
    s2 = s;
}
```

Here, both `s1` and `s2` point at the same string and a naive implementation of `save_state` could end up by copying the string twice. From a `save_state` perspective this is simply inefficient. However, it makes `restore_state` incorrect since it would unpack the two strings into different areas of memory destroying the original aliasing information. The current version of *TxCore* will pack and unpack separate object references.

Direct use of StateManager

The examples throughout this manual have always derived user classes from `LockManager`. The reasons for this are twofold. Firstly, and most importantly, the serialisability constraints of atomic actions require it, and secondly it reduces the need for programmer intervention. However, if only access to *TxCore's* persistence and recovery mechanisms is required, direct derivation of a user class from `StateManager` is possible.

Classes derived directly from `StateManager` must make use of its state management mechanisms explicitly (these interactions are normally undertaken by `LockManager`). From a programmer's point of view this amounts to making appropriate use of the operations `activate`, `deactivate` and `modified`, since `StateManager's` constructors are effectively identical to those of `LockManager`.

```
boolean activate ()
boolean activate (String storeRoot)
```

`Activate` loads an object from the object store. The object's UID must already have been set via the constructor and the object must exist in the store. If the object is successfully read then `restore_state` is called to build the object in memory. `Activate` is idempotent so that once an object has been activated further calls are ignored. The parameter represents the root name of the object store to search for the object. A value of null means use the default store.

```
boolean deactivate ()
boolean deactivate (String storeRoot)
```

The inverse of `activate`. First calls `save_state` to build the compacted image of the object which is then saved in the object store. Objects are only saved if they have been modified since they were activated. The parameter represents the root name of the object store into which the object should be saved. A value of null means use the default store.

```
void modified ()
```

Must be called prior to modifying the object in memory. If it is not called the object *will not* be saved in the object store by `deactivate`.

Tools

Introduction

This chapter explains how to start and use the tools framework and what tools are available.

Starting the Transaction Service tools

The way to start the transaction service tools differs on the operating system being used:

Windows:

Double click on the ‘Start Tools’ link in the JBoss Transaction Service program group in the start menu.

UNIX:

Start a bash shell and type:

```
cd < JBossTS INSTALL DIRECTORY >
./run-tools.sh
```

Once you have done this the tools window will appear. This is the launch area for all of the tools shipped with the JBoss Transaction Service. At the top of the window you will notice a menu bar (see Figure 7).



This menu bar has four menus:

The **F**ile menu:



Open **JMX** Browser – this displays the JMX browser window (see **Using the JMX Browser** for more information on how to use the JMX browser).

Open **O**bject Store Browser – this displays the JBossTS Object Store browser window (see **Using the Object Store Browser** for more information on how to use the Object Store browser).

Settings – this option opens the settings dialog which lets you configure the different tools available.

Exit – this closes the tools window and exits the application, any unsaved/unconfirmed changes will be lost.

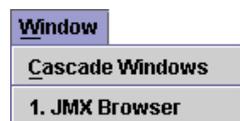
The **P**erformance menu:



Open – this opens a performance window – see the section named ‘Using the Performance Tool’ for more information on the performance tool.

Close All – this closes all of the currently open performance windows – see the section named ‘Using the Performance Tool’ for more information on the performance tool.

The **W**indow menu:



Cascade Windows – this arranges the windows in a diagonal line to you find a specific window.

1. xxxxxx – For each window currently visible an extra menu option will be available here. Selecting this menu option will bring the associated window to the front of the desktop.

The **H**elp menu:



About – this displays the about window which displays the product information.

Using the Performance Tool

The performance tool can be used to display performance information about the transaction service. This information is gathered using the Performance JMX bean which means that the transaction service needs to be integrated into an Application Server to give any performance information.

The performance information is displayed via a multi-series graph. To view this graph simply open a performance window by selecting Performance > Open (see Figure 8).

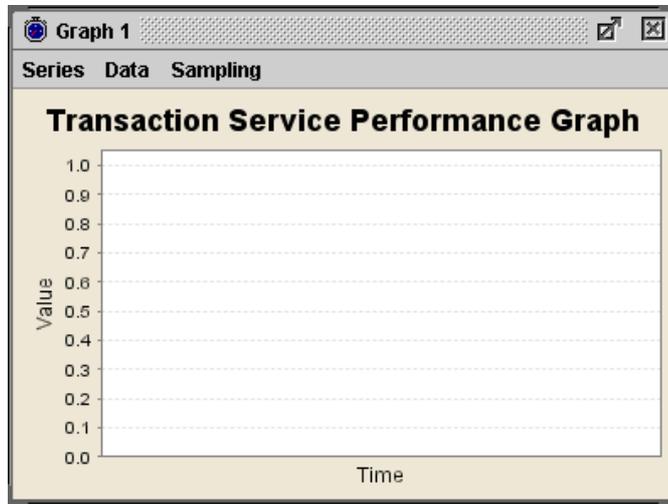


Figure 8 - Performance window

This window contains a multi-series graph which can display the following information:

- Number of transactions.
- Number of committed transactions.
- Number of aborted transactions.
- Number of nested transactions.
- Number of heuristics raised.

To turn these series on and off simply select the menu option from the series menu:



When series are turned on they appear in the legend at the bottom of the graph. The colour next to the series name (e.g. Transactions Created) is the colour of the line representing that data.



The data shown is graphed against time. The Y-axis represents the number of transactions and the X-axis represents time.

At any point the sampling of data can be stopped and restarted using the ‘Sampling’ menu and the data currently visible in the graph can be saved to a Comma Separate Values (CSV) file for importing the data into a spreadsheet application using the ‘Save to .csv’ menu option from the ‘Data’ menu.

Using the JMX Browser

To open the JMX browser window click on the **File** menu and then click the **Open JMX Browser** option. The JMX browser window will then be displayed (see **Figure 9**).

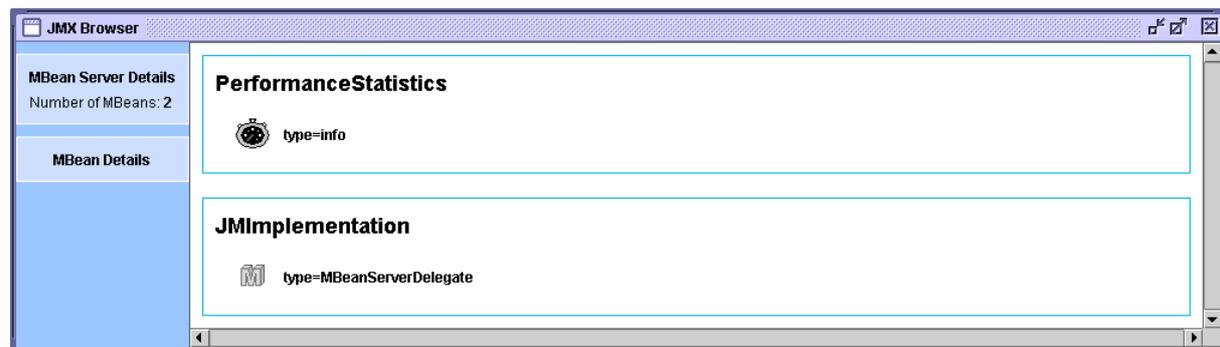


Figure 9 - JMX Browser window.

The window is made up of two main sections: the details panel and the MBean panel. The MBean panel displays the MBeans exposed by the MBean server. These are grouped by domain name. The details panel displays information about the currently selected MBean. To select an MBean just left-click it with the mouse and it will become highlighted. The information displayed in the details panel is as follows (see **Figure 10** for an example):

- The total number of MBeans registered on this server,
- The number of constructors exposed by this MBean,
- The number of attributes exposed by this MBean,
- The number of operations exposed by this MBean,
- The number of notifications exposed by this MBean,
- A brief description of the MBean.

There is also a **View** link which when clicked displays the attributes and operations exposed by this MBean. From there you can view readable attributes, alter writeable attributes and invoke operations.

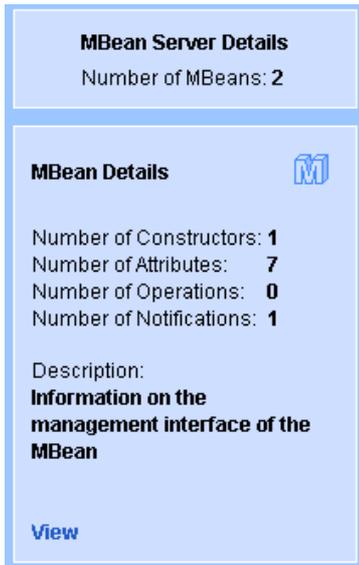


Figure 10 - An example of what the details panel displays.

Using Attributes and Operations

When the **View** link is clicked the View JMX Attributes and Operations window is displayed (see **Figure 11**). From here you can view all readable attributes exposed by the selected MBean. You can also alter writeable attributes. If an attribute is read-only then you will not be able to alter an attributes value. To alter an attributes value just double click on the current value and enter the new value. If the button is enabled then you can click this to view a more suitable editing method. If the attribute type is a JMX object name then clicking this button will display the JMX attributes and operations for that object.

At any point you can click the button to refresh the attribute values. If an exception occurs while retrieving the value of an attribute the exception will be displayed in place of the attributes value.

You can also invoke operations upon an MBean. A list of operations exposed by an MBean is displayed below the attributes list. To invoke an operation simply select it from the list and click the button. If the operation requires parameters a further window will be displayed, from this window you must specify values for each of the parameters required (see **Figure 12**). You specify parameter values in the same way as you specify JMX attribute values. Once you have specified a value for each of the parameters click the **Invoke** button to perform the invocation.

Once the method invocation has completed its return value will be displayed.

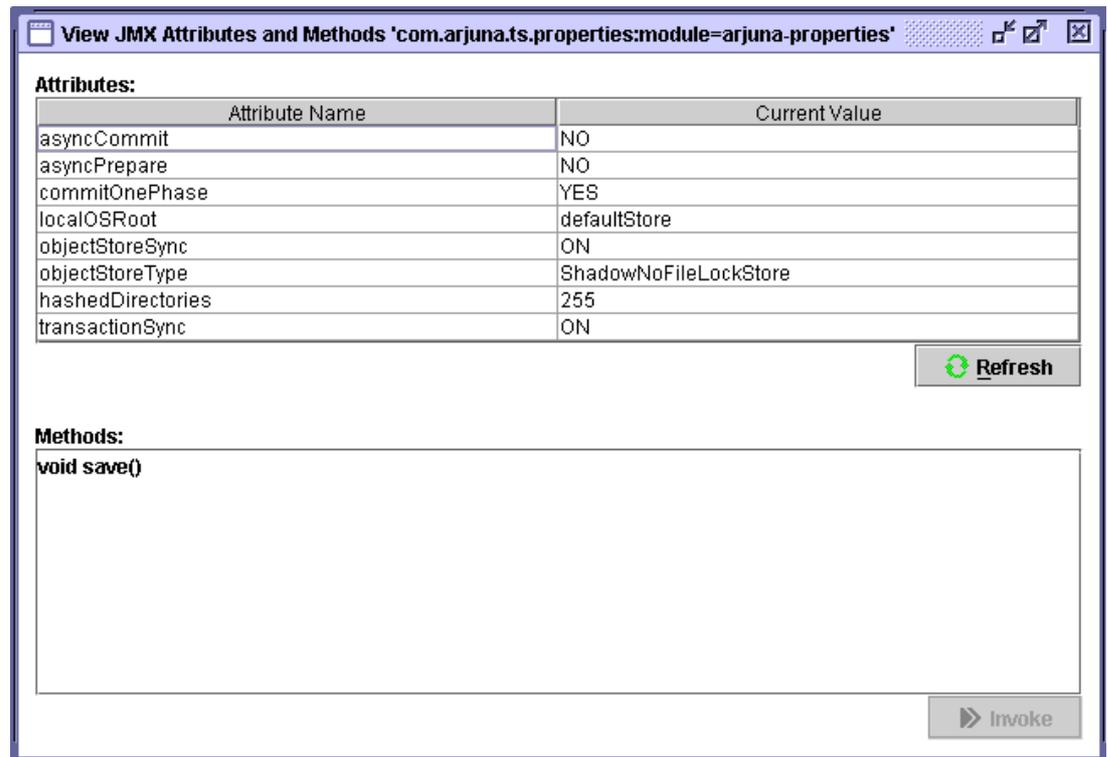


Figure 11 - View JMX Attributes and Operations window.

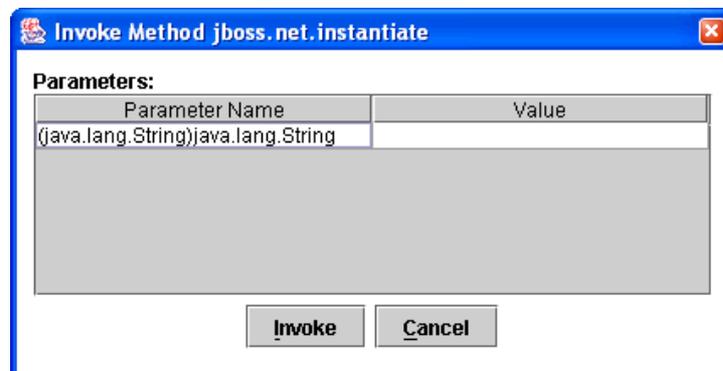


Figure 12 - Invoke Operation Parameters.

Using the Object Store Browser

To open the Object Store browser window click on the **File** menu and then click the **Open Object Store Browser** option. The Object Store browser window will then be displayed (see Figure 13).

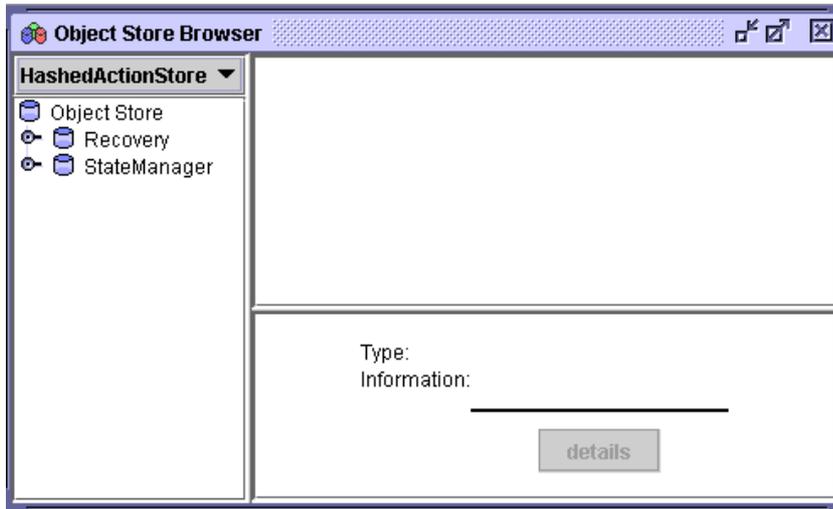


Figure 13 - Object Store Browser window

The object store browser window is split into four sections:

Object Store Roots	Objects
Object Store Hierarchy	Object Details

- Object Store Roots – this is a pull down of the currently available object store roots. Selecting an option from the list will repopulate the hierarchy view with the contents of the selected root.
- Object Store Hierarchy – this is a tree which shows the current object store hierarchy. Selecting a node from this tree will display the objects stored in that location.

- Objects – this is a list of icons which represent the objects stored in the selected location.
- Object Details – this shows information about the currently selected object (only if the object's type is known to the state viewer repository see **Writing an OSV** for information on how to write a object state viewers).

Object State Viewers (OSV)

When an object is selected in the objects pane of the main window the registered Object State Viewer (or OSV) for that object type is invoked. An OSV's job is to make information available via the user interface to the user to show information about the selected object. Distributed with the standard tools is an OSV for Atomic Actions, the OSV displays information on the Abstract Records in it's various lists (e.g. heuristic, failed, read-only, etc). It is also possible to write your own OSVs which can be used to display information about object types you have defined. This subject is covered next.

Writing an OSV

Writing an OSV plugin allows you to extend the capabilities of the Object Store browser to show the state of user defined abstract records. An OSV plug-in is simply a class which implements the interface:

```
com.arjuna.ats.tools.objectstorebrowser.stateviewers.StateViewerInterface
```

It must be packaged in a JAR within the plugins directory. This example shows how to create an OSV plugin for an abstract record subclass which looks as follows:

```
public class SimpleRecord extends AbstractRecord
{
    private int _value = 0;

    .....

    public void increase()
    {
        _value++;
    }

    public int get()
    {
        return _value;
    }

    public String type()
    {
        return "/StateManager/AbstractRecord/SimpleRecord";
    }

    public boolean restore_state(InputObjectState os, int i)
    {
        boolean returnValue = true;

        try
        {
```

```
        _value = os.unpackInt();
    }
    catch (java.io.IOException e)
    {
        returnValue = false;
    }

    return returnValue;
}

public boolean save_state(OutputObjectState os, int i)
{
    boolean returnValue = true;

    try
    {
        os.packInt(_value);
    }
    catch (java.io.IOException e)
    {
        returnValue = false;
    }

    return returnValue;
}
}
```

When this abstract record is viewed in the object store browser it would be nice to see the current value. This is easy to do as we can read the state into an instance of our abstract record and call `getValue()`. The following is the object store browser plug-in source code:

```
public class SimpleRecordOSVPlugin implements StateViewerInterface
{
    /**
     * A uid node of the type this viewer is registered against has been
    expanded.
     * @param os
     * @param type
     * @param manipulator
     * @param node
     * @throws ObjectStoreException
     */
    public void uidNodeExpanded(ObjectStore os,
                               String type,
                               ObjectStoreBrowserTreeManipulationInterface
                                   manipulator,
                               UidNode node,
                               StatePanel infoPanel)
        throws ObjectStoreException
    {
        // Do nothing
    }

    /**
     * An entry has been selected of the type this viewer is registered
    against.
     *
     * @param os
     * @param type
     * @param uid
     * @param entry
     * @param statePanel
     */
}
```

```

        * @throws ObjectStoreException
        */
    public void entrySelected(ObjectStore os,
                             String type,
                             Uid uid,
                             ObjectStoreViewEntry entry,
                             StatePanel statePanel)
        throws ObjectStoreException
    {
        SimpleRecord rec = new SimpleRecord();

        if ( rec.restore_state( os.read_committed(uid, type),
ObjectTypes.ANDPERSISTENT ) )
        {
            statePanel.setData( "Value", rec.getValue() );
        }
    }

    /**
     * Get the type this state viewer is intended to be registered against.
     * @return
     */
    public String getType()
    {
        return "/StateManager/AbstractRecord/SimpleRecord";
    }
}

```

The method `uidNodeExpanded` is invoked when a UID (Unique Identification) representing the given type is expanded in the object store hierarchy tree. This is not required by this plugin as this abstract record is not visible in the object store directly it is only viewable via one of the lists in an atomic action. The method `entrySelected` is invoked when an entry is selected from the object view which represents an object with the given type. In both methods the `StatePanel` is used to display information regarding the state of the object. The state panel has the following methods that assist in display this information:

- `setInfo(String info)` – This method can be used to show general information.
- `setData(String name, String value)` – This method is used to put information into the table which is displayed by the object store browser tool.
- `enableDetailsButton(DetailsButtonListener listener)` – This method is used to enable the details button. The listener interface allows a plug-in to be informed when the button is pressed. It is up to the plug-in developer to decide how to display this further information.

In this example we read the state from the object store and use the value returned by `getValue()` to put an entry into the state panel table. The `getType()` method returns the type this plug-in is to be registered against.

To add this plug-in to the object store browser it is necessary to package it into a JAR (Java Archive) file with a name that is prefixed with 'osbv-'. The JAR file must contain certain information within the manifest file so that the object store browser knows which classes are plug-ins. All of this can be performed using an Apache ANT (<http://ant.apache.org>) script, as follows:

```
<jar jarfile="osbv-simplerecord.jar">
  <fileset dir="build" includes="*.class"/>
  <manifest>
    <section name="arjuna-tools-objectstorebrowser">
      <attribute name="plugin-classname-1" value="
SimpleRecordOSVPlugin "/>
    </section>
  </manifest>
</jar>
```

Once the JAR has been created with the correct information in the manifest file it just needs to be placed in the 'bin/tools/plugins' directory.

Constructing a Transactional Objects for Java application

Application construction

There are two distinct phases to the development of a *TxCore* application:

- Developing new classes with certain characteristics (for example, Persistent, Recoverable, Concurrency Controlled).
- Developing the application(s) that make use of the new classes of objects.

Although these two phases may be performed in parallel and by a single person, we shall refer to the first step as the job of the *class developer* and the second as the job of the *applications developer*. The class developer will be concerned about defining appropriate `save_state` and `restore_state` operations for the class, setting appropriate locks in operations, and invoking the appropriate *TxCore* class constructors. The applications developer will be more concerned with defining the general structure of the application, particularly with regard to the use of atomic actions.

This chapter illustrates the points made in previous sections by outlining a simple application: in this case a simple FIFO `Queue` class for integer values will be developed. The implementation of the `Queue` will be with a doubly linked list structure, and it will be implemented as a single object. We shall be using this example throughout the rest of this manual to help illustrate the various mechanisms provided by *TxCore*. While this is an unrealistic example application it enables all of the *TxCore* modifications to be described without requiring in depth knowledge of the application code.

In the rest of this chapter we shall assume that the application is not distributed. If this is not the case, then context information must be propagated either implicitly or explicitly.

Queue description

The queue is a traditional FIFO queue, where elements are added to the front and removed from the back. The operations provided by the queue class allow the values to be placed on to the queue (`enqueue`) and to be removed from it (`dequeue`), and it is also possible to change or inspect the values of elements in the queue. In this example implementation, an array is

used to represent the queue. A limit of `QUEUE_SIZE` elements has been imposed for this example.

The Java interface definition of this simple queue class is given below:

```
public class TransactionalQueue extends LockManager
{
    public TransactionalQueue (Uid uid);
    public TransactionalQueue ();
    public void finalize ();

    public void enqueue (int v) throws OverFlow, UnderFlow,
                               QueueError, Conflict;
    public int dequeue  () throws OverFlow, UnderFlow,
                               QueueError, Conflict;

    public int queueSize ();
    public int inspectValue (int i) throws OverFlow,
                               UnderFlow, QueueError, Conflict;
    public void setValue (int i, int v) throws OverFlow,
                               UnderFlow, QueueError, Conflict;

    public boolean save_state (OutputObjectState os, int ObjectType);
    public boolean restore_state (InputObjectState os, int ObjectType);
    public String type ();

    public static final int QUEUE_SIZE = 40; // maximum size of the queue

    private int[QUEUE_SIZE] elements;
    private int numberOfElements;
};
```

Constructors and destructors

As stated in the previous section, to use an existing persistent object requires the use of a special constructor that is required to take the `Uid` of the persistent object; the implementation of such a constructor is given below:

```
public TransactionalQueue (Uid u)
{
    super(u);

    numberOfElements = 0;
}
```

The constructor that creates a new persistent object is similar:

```
public TransactionalQueue ()
{
    super(ObjectType.ANDPERSISTENT);

    numberOfElements = 0;

    try
    {
        AtomicAction A = new AtomicAction();

        A.begin(0); // Try to start atomic action
    }
}
```

```
        // Try to set lock

        if (setlock(new Lock(LockMode.WRITE), 0) == LockResult.GRANTED)
        {
            A.commit(true); // Commit
        }
        else // Lock refused so abort the atomic action
            A.rollback();
        }
        catch (Exception e)
        {
            System.err.println("Object construction error: "+e);
            System.exit(1);
        }
    }
}
```

The use of an atomic action within the constructor for a new object follows the guidelines outlined earlier and ensures that the object's state will be written to the object store when the appropriate top level atomic action commits (which will either be the action A or some enclosing action active when the `TransactionalQueue` was constructed). The use of atomic actions in a constructor is simple: an action must first be declared and its `begin` operation invoked; the operation must then set an appropriate lock on the object (in this case a `WRITE` lock must be acquired), then the main body of the constructor is executed. If this is successful the atomic action can be committed, otherwise it is aborted.

The destructor of the queue class is only required to call the `terminate` operation of `LockManager`

```
public void finalize ()
{
    super.terminate();
}
```

save_state, restore_state and type

The implementations of `save_state` and `restore_state` are relatively simple for this example:

```
public boolean save_state (OutputObjectState os, int ObjectType)
{
    if (!super.save_state(os, ObjectType))
        return false;

    try
    {
        os.packInt(numberOfElements);

        if (numberOfElements > 0)
        {
            for (int i = 0; i < numberOfElements; i++)
                os.packInt(elements[i]);
        }

        return true;
    }
    catch (IOException e)
```

```
    {
        return false;
    }
}
public boolean restore_state (InputObjectState os, int ObjectType)
{
    if (!super.restore_state(os, ObjectType))
        return false;

    try
    {
        numberOfElements = os.unpackInt();

        if (numberOfElements > 0)
        {
            for (int i = 0; i < numberOfElements; i++)
                elements[i] = os.unpackInt();
        }

        return true;
    }
    catch (IOException e)
    {
        return false;
    }
}
```

Because the `Queue` class is derived from the `LockManager` class, the operation type should be:

```
public String type ()
{
    return "/StateManager/LockManager/TransactionalQueue";
}
```

enqueue/dequeue operations

If the operations of the queue class are to be coded as atomic actions, then the `enqueue` operation could have the structure given below (the `dequeue` operation would be similarly structured):

```
public void enqueue (int v) throws OverFlow, UnderFlow, QueueError
{
    AtomicAction A = new AtomicAction();
    boolean res = false;

    try
    {
        A.begin(0);

        if (setlock(new Lock(LockMode.WRITE), 0) == LockResult.GRANTED)
        {
            if (numberOfElements < QUEUE_SIZE)
            {
                elements[numberOfElements] = v;
                numberOfElements++;
                res = true;
            }
            else

```

```
        {
            A.rollback();
            throw new UnderFlow();
        }
    }

    if (res)
        A.commit(true);
    else
    {
        A.rollback();
        throw new Conflict();
    }
}
catch (Exception e1)
{
    throw new QueueError();
}
}
```

queueSize

The implementation of `queueSize` is shown below:

```
public int queueSize () throws QueueError, Conflict
{
    AtomicAction A = new AtomicAction();
    int size = -1;

    try
    {
        A.begin(0);

        if (setlock(new Lock(LockMode.READ), 0) == LockResult.GRANTED)
            size = numberOfElements;

        if (size != -1)
            A.commit(true);
        else
        {
            A.rollback();

            throw new Conflict();
        }
    }
    catch (Exception e1)
    {
        throw new QueueError();
    }

    return size;
}
```

inspectValue/setValue operations

The implementation of `inspectValue` is shown below. `setValue` is similar, and not shown.

```
public int inspectValue (int index) throws UnderFlow,
                               OverFlow, Conflict, QueueError
{
    AtomicAction A = new AtomicAction();
    boolean res = false;
    int val = -1;

    try
    {
        A.begin();

        if (setlock(new Lock(LockMode.READ), 0) == LockResult.GRANTED)
        {
            if (index < 0)
            {
                A.rollback();
                throw new UnderFlow();
            }
            else
            {
                // array is 0 - numberOfElements -1

                if (index > numberOfElements -1)
                {
                    A.rollback();
                    throw new OverFlow();
                }
                else
                {
                    val = elements[index];
                    res = true;
                }
            }
        }

        if (res)
            A.commit(true);
        else
        {
            A.rollback();
            throw new Conflict();
        }
    }
    catch (Exception e1)
    {
        throw new QueueError();
    }

    return val;
}
```

The client

Rather than show all of the code for the client, we shall concentrate on a representative portion. Before invoking operations on the object, the client must obviously first bind to it. In the local case this simply requires the client to create an instance of the object.

```
public static void main (String[] args)
{
    TransactionalQueue myQueue = new TransactionalQueue();
}
```

Before invoking one of the queue's operations, the client starts a transaction. The `queueSize` operation is shown below:

```
AtomicAction A = new AtomicAction();
int size = 0;

try
{
    A.begin(0);

    try
    {
        size = queue.queueSize();
    }
    catch (Exception e)
    {
    }

    if (size >= 0)
    {
        A.commit(true);

        System.out.println("Size of queue: "+size);
    }
    else
        A.rollback();
}
catch (Exception e)
{
    System.err.println("Caught unexpected exception!");
}
```

Comments

Since the queue object is persistent, then the state of the object will survive any failures of the node on which it is located. The state of the object that will survive is that produced by the last top-level committed atomic action performed on the object. If it is the intention of an application to perform two `enqueue` operations atomically, for example, then this can be done by nesting the `enqueue` operations in another enclosing atomic action. In addition, concurrent operations on such a persistent object will be serialised, thereby preventing inconsistencies in the state of the object. However, since the elements of the queue objects are not individually concurrency controlled, certain combinations of concurrent operation invocations will be executed serially, whereas logically they could be executed concurrently. For example, modifying the states of two different elements in the queue. In the next section we address some of these issues.

Configuration options

Options

The following table shows the configuration features, with default values shown in italics. More details about each option can be found in the relevant sections of this document.

Configuration Name	Possible Values	Description
com.arjuna.ats.arjuna.objectstore.storeSync	<i>ON/OFF</i>	Turns synchronization of the object store on or off. Use with caution.
com.arjuna.ats.arjuna.objectstore.storeType	<i>ShadowStore/ShadowNoFileLockStore/JDBCStore/HashedStore</i>	Specify the type of object store implementation to use.
com.arjuna.ats.arjuna.objectstore.hashedDirectories	<i>255/any integer value</i>	Set the number of directories to hash object states over for the HashedStore object store implementation.
com.arjuna.ats.txoj.lockstore.lockStoreType	<i>BasicLockStore/BasicPersistentLockStore</i>	Specify the type of the lock store implementation to use.
com.arjuna.ats.txoj.lockstore.lockStoreDir	Windows: <i>.\LockStore</i> Unix: <i>./LockStore</i>	Specify the location of the lock store.
com.arjuna.ats.arjuna.objectstore.objectStoreDir	<i>Any location the application can write to.</i>	Specify the location of the object store.
com.arjuna.ats.arjuna.objectstore.localOSRoot	<i>defaultStore</i>	Specify the name of the object store root.
com.arjuna.ats.arjuna.coordinator.actionStore	<i>ActionStore/HashedActionStore/JDBCActionStore</i>	The transaction log implementation to use.
com.arjuna.ats.arjuna.coordinator.asyncCommit	<i>YES/NO</i>	Turns on or off (default) asynchronous commit.
com.arjuna.ats.arjuna.coordinator.asyncPrepare	<i>YES/NO</i>	Turns on or off (default) asynchronous prepare.
com.arjuna.ats.arjuna.objectstore.transactionSync	<i>ON/OFF</i>	Turns synchronization of the object store on or off. Use with caution.
com.arjuna.ats.arjuna.objectstore.jdbcUserDbAccess	<i>JDBCAccess class name</i>	The JDBCAccess implementation to use for user-level object stores.

<code>com.arjuna.ats.arjuna.objectstore.jdbcTxDbAccess</code>	JDBCAccess class name	The JDBCAccess implementation to use for transaction object stores.
<code>com.arjuna.ats.arjuna.coordinator.commitOnePhase</code>	YES/NO	Enable or disable the one-phase commit optimization.
<code>com.arjuna.ats.arjuna.coordinator.readonlyOptimisation</code>	YES/NO	Enable or disable read-only optimization for the second phase abort.
<code>com.arjuna.ats.arjuna.coordinator.enableStatistics</code>	YES/NO	Start/stop collecting transaction statistic information.
<code>com.arjuna.ats.arjuna.coordinator.startDisabled</code>	YES/NO	Start with the transaction system enabled or disabled. Toggle via the <code>com.arjuna.ats.arjuna.coordinator.TxControl</code> class.
<code>com.arjuna.ats.arjuna.coordinator.defaultTimeout</code>	Integer	Timeout in milliseconds

Table 2 TxCore configuration options.

Object store implementations

The ObjectStore

In this appendix we shall examine the various *TxCore* object store implementations and give guidelines as to how other implementations may be created and plugged into an application.

This release of *JBossTS* contains several different implementations of a basic object store. Each serves a particular purpose and is generally optimised for that purpose. All of the implementations are derived from the `ObjectStore` interface. This defines the minimum operations which must be provided in order for an object store implementation to be used by *JBossTS*. The default object store implementation can be overridden at runtime by setting the `com.arjuna.ats.arjuna.objectstore.objectStoreType` property variable to one of the types described below.

```
/*
 * This is the base class from which all object store types are derived.
 * Note that because object store instances are stateless, to improve
 * efficiency we try to only create one instance of each type per process.
 * Therefore, the create and destroy methods are used instead of new
 * and delete. If an object store is accessed via create it *must* be
 * deleted using destroy. Of course it is still possible to make use of
 * new and delete directly and to create instances on the stack.
 */

public class ObjectStore
{
public static final int OS_COMMITTED;
public static final int OS_COMMITTED_HIDDEN;
public static final int OS_HIDDEN;
public static final int OS_INVISIBLE;
public static final int OS_ORIGINAL;
public static final int OS_SHADOW;
public static final int OS_UNCOMMITTED;
public static final int OS_UNCOMMITTED_HIDDEN;
public static final int OS_UNKNOWN;
public ObjectStore (ClassName type);
public ObjectStore (ClassName type, String osRoot);
public ObjectStore (String osRoot);
public synchronized boolean allObjUids (String s, InputObjectState buff)
    throws ObjectStoreException;
public synchronized boolean allObjUids (String s, InputObjectState buff,
    int m) throws ObjectStoreException;

public synchronized boolean allTypes (InputObjectState buff)
    throws ObjectStoreException;
public synchronized int currentState(Uid u, String tn)
    throws ObjectStoreException;
}
```

```

public synchronized boolean commit_state (UId u, String tn)
    throws ObjectStoreException;
public synchronized boolean hide_state (UId u, String tn)
    throws ObjectStoreException;
public synchronized boolean reveal_state (UId u, String tn)
    throws ObjectStoreException;
public synchronized InputObjectState read_committed (UId u, String tn)
    throws ObjectStoreException;
public synchronized InputObjectState read_uncommitted (UId u, String tn)
    throws ObjectStoreException;
public synchronized boolean remove_committed (UId u, String tn)
    throws ObjectStoreException;
public synchronized boolean remove_uncommitted (UId u, String tn)
    throws ObjectStoreException;
public synchronized boolean write_committed (UId u, String tn,
    OutputObjectState buff)
    throws
ObjectStoreException;
public synchronized boolean write_uncommitted (UId u, String tn,
    OutputObjectState buff)
    throws
ObjectStoreException;
public static void printState (PrintStream strm, int res);
};

```

JBossTS programmers need not usually interact with any of the object store implementations directly other than possibly to create them in the first place (even this is not necessary if the default store type is used as *JBossTS* will create stores as necessary). All stores manipulate instances of the class `ObjectState` which are named using a type (via the object's `type()` operation) and a `UId`. For atomic actions purposes object states in the store can be principally in two distinct states: `OS_COMMITTED`, and `OS_UNCOMMITTED`. An object state starts in the `OS_COMMITTED` state but when modified under the control of an atomic action a new *second* object state may be written that is in the `OS_UNCOMMITTED` state. If the action commits this second object state replaces the original and becomes `OS_COMMITTED`. If the action aborts, this second object state is simply discarded. All of the implementations provided with this release handle these state transitions by making use of *shadow* copies of object states, however, any other implementation that maintains this abstraction is permissible. Object states may become *hidden* (and thus inaccessible) under the control of the crash recovery system.

Browsing of the contents of a store is possible through the `allTypes` and `allObjUids` operations. `allTypes` returns an `InputObjectState` containing all of the type names of all objects in a store, terminated by a null name. `allObjUids` returns an `InputObjectState` that contains all of the `Uids` of all objects of a given type terminated by the special `UId.nullUId()`.

Persistent object stores

This section briefly describes the characteristics and optimisations of each of the supplied implementations of the persistent object store. Persistent object states are mapped onto the structure of the file system supported by the host operating system.

Common functionality

In addition to the features mentioned earlier all of the supplied persistent object stores obey the following rules:

- Each object state is stored in its own file that is named using the `UId` of the object.
- The type of an object (as given by the `type()` operation) determines the directory into which the object is placed.
- All of the stores have a *common* root directory that is determined when *JBossTS* is configured. This directory name is automatically prepended to any store specific root information.
- All stores also have the notion of a localised root directory that is automatically prepended to the type of the object to determine the ultimate directory name. The localised root name is specified when the store is created. By default the localised root name is `defaultStore`.

Thus the default directory layout is then:

```
<ObjectStore root Directory from configure>
/JBossTS/ObjectStore/
<ObjectStore Type1>                                     FragmentedStore/
  <Default root>                                       defaultStore/
    <StateManager>                                     StateManager
      <LockManager>                                   LockManager/
        <User Types>
  <Localised root 2>                                   myStore/
    <StateManager>                                   StateManager/

<ObjectStore Type2>                                     ActionStore/
  <Default root>                                       defaultStore/
```

The shadowing store

This is the original version of the object store as provided in prior releases and is implemented by the class `ShadowingStore`. It is simple but slow. It uses pairs of files to represent objects (the shadow version and the committed version) and files are opened, locked, operated upon, unlocked and closed on every interaction with the object store. Thus significant portions of time can be spent in the system simply opening, closing and renaming files, all of which are very expensive operations.

If overriding the object store implementation, the type of this object store is “ShadowingStore”.

No file-level locking

Since transactional objects are concurrency controlled through `LockManager`, it is not necessary to impose additional locking at the file level, as the basic `ShadowingStore` implementation does. Therefore, the default object store implementation for *JBossTS*,

`ShadowNoFileLockStore`, relies upon user-level locking. This enables it to provide better performance than the `ShadowingStore` implementation.

If overriding the object store implementation, the type of this object store is “`ShadowNoFileLockStore`”.

The hashed store

The `HashedStore` has the same structure for object states as the shadowing stores but has an alternate directory structure that is better suited to storing large numbers of objects of the same type. Using this store objects are scattered amongst a set of directories by applying a hashing function to the object's `Uid`. By default 255 sub-directories are used. However, this can be overridden by setting the `HASHED_DIRECTORIES` environment variable accordingly.

If overriding the object store implementation, the type of this object store is “`HashedStore`”.

The JDBC store

The `JDBCStore` uses a JDBC database to save persistent object states; when used in conjunction with the *Transactional Objects for Java* API nested transaction support is available. In the current implementation, all object states are stored as Binary Large Objects (BLOBs) within the same table. The limitation on object state size imposed by using BLOBs is 64k; if an attempt is made to store an object state which exceeds this limit an error will be output and the state will not be stored. The transaction will subsequently be forced to roll back.

When using the JDBC object store, the application must provide an implementation of the following interface, located in the `com.arjuna.ats.arjuna.objectstore` package:

```
public interface JDBCAccess
{
    public Connection getConnection () throws SQLException;
    public void putConnection (Connection conn) throws SQLException;
    public void initialise (ObjectName objName);
}
```

The implementation of this class is responsible for providing the `Connection` which the `JDBC ObjectStore` will use to save and restore object states:

- `getConnection`: returns the `Connection` to use. This method will be called whenever a connection is required and the implementation should use whatever policy is necessary for determining what connection to return. This method need not return the same `Connection` instance more than once.
- `putConnection`: this method will be called to return one of the `Connections` acquired from `getConnection`. `Connections` are returned if any errors occur when using them.
- `initialise`: this can be used to pass additional arbitrary information to the implementation.

The JDBC object store will initially request the number of Connections defined in the `com.arjuna.ats.arjuna.objectstore.jdbcPoolSizeInitial` property and will use no more than defined in the `com.arjuna.ats.arjuna.objectstore.jdbcPoolSizeMaximum` property.

The implementation of the `JDBCAccess` interface to use should be set in the `com.arjuna.ats.arjuna.objectstore.jdbcUserDbAccess` property variable.

If overriding the object store implementation, the type of this object store is “JDBCStore”.

A JDBC object store can be used for managing the transaction log. In this case, the transaction log implementation should be set to “JDBCActionStore” and the `JDBCAccess` implementation must be provided via the `com.arjuna.ats.arjuna.objectstore.jdbcTxDbAccess` property variable. In this case, the default table name is *JBossTSTxTable*.

Note: It is possible to use the same `JDBCAccess` implementation for both the user object store and also the transaction log.

The cached store

This object store used the hashed object store, but does not read or write states to the persistent backing store immediately. It maintains the states in a volatile memory cache and either flushes the cache periodically or when it is full. The failure semantics associated with this object store are different to the normal persistent object stores, because a failure could result in states in the cache being lost.

If overriding the object store implementation, the type of this object store is “CachedStore”.

The store can be configured with the following properties:

- `com.arjuna.ats.internal.arjuna.objectstore.cacheStore.hash` sets the number of internal stores to hash the states over. The default value is 128.
- `com.arjuna.ats.internal.arjuna.objectstore.cacheStore.size` is the maximum size the cache can reach before a flush is triggered. The default is 10240 bytes.
- `com.arjuna.ats.internal.arjuna.objectstore.cacheStore.removedItems` is the maximum number of removed items that the cache can contain before a flush is triggered. By default, calls to remove a state that is in the cache will simply remove the state from the cache, but leave a blank entry (rather than remove the entry immediately, which would affect the performance of the cache). When triggered, these entries are removed from the cache. The default value is twice the size of the hash.
- `com.arjuna.ats.internal.arjuna.objectstore.cacheStore.workItems` is the maximum number of items that are allowed to build up in the cache before it is flushed. The default value is 100.

- `com.arjuna.ats.internal.arjuna.objectstore.cacheStore.scanPeriod` sets the time in milliseconds for periodically flushing the cache. The default is 120 seconds.
- `com.arjuna.ats.internal.arjuna.objectstore.cacheStore.sync` determines whether flushes of the cache are sync-ed to disk. The default is OFF. To enable, set to ON.

Class definitions

Introduction

This appendix contains an overview of those classes that the application programmer will typically use. The aim of this appendix is to provide a quick reference guide to these classes for use when writing applications in *TxCore*. For clarity only the public and protected interfaces of the classes will be given.

Class library

LockManager

```
public class LockResult
{
    public static final int GRANTED;
    public static final int REFUSED;
    public static final int RELEASED;
};

public class ConflictType
{
    public static final int CONFLICT;
    public static final int COMPATIBLE;
    public static final int PRESENT;
};

public abstract class LockManager extends StateManager
{
    public static final int defaultRetry;
    public static final int defaultTimeout;
    public static final int waitTotalTimeout;

    public final synchronized boolean releaselock (UId lockUId);
    public final synchronized int setlock (Lock toSet);
    public final synchronized int setlock (Lock toSet, int retry);
    public final synchronized int setlock (Lock toSet, int retry, int
sleepTime);
    public void print (PrintStream strm);
    public String type ();
    public boolean save_state (OutputObjectState os, int ObjectType);
    public boolean restore_state (InputObjectState os, int ObjectType);

    protected LockManager ();
    protected LockManager (int ot);
    protected LockManager (int ot, ObjectName attr);
    protected LockManager (UId storeUId);
    protected LockManager (UId storeUId, int ot);
    protected LockManager (UId storeUId, int ot, ObjectName attr);
```

```
protected void terminate ();
};
```

StateManager

```
public class ObjectStatus
{
public static final int PASSIVE;
public static final int PASSIVE_NEW;
public static final int ACTIVE;
public static final int ACTIVE_NEW;
};

public class ObjectType
{
public static final int RECOVERABLE;
public static final int ANDPERSISTENT;
public static final int NEITHER;
};

public abstract class StateManager
{
public boolean restore_state (InputObjectState os, int ot);
public boolean save_state (OutputObjectState os, int ot);
public String type ();

public synchronized boolean activate ();
public synchronized boolean activate (String rootName);
public synchronized boolean deactivate ();
public synchronized boolean deactivate (String rootName);
public synchronized boolean deactivate (String rootName, boolean commit);

public synchronized int status ();
public final Uid get_uid ();
public void destroy ();
public void print (PrintStream strm);

protected void terminate ();

protected StateManager ();
protected StateManager (int ot);
protected StateManager (int ot, ObjectName objName);
protected StateManager (Uid objUid);
protected StateManager (Uid objUid, int ot);
protected StateManager (Uid objUid, int ot, ObjectName objName);
protected synchronized final void modified ();
};
```

Input/OutputObjectState

```
class OutputObjectState extends OutputBuffer
{
public OutputObjectState (Uid newUid, String typeName);

public boolean notempty ();
public int size ();
public Uid stateUid ();
public String type ();
};
```

```
class InputObjectState extends ObjectState
{
public OutputObjectState (Uid newUid, String typeName, byte[] b);

public boolean notempty ();
public int size ();
public Uid stateUid ();
public String type ();
};
```

Input/OutputBuffer

```
public class OutputBuffer
{
public OutputBuffer ();

public final synchronized boolean valid ();
public synchronized byte[] buffer();
public synchronized int length ();

    /* pack operations for standard Java types */

public synchronized void packByte (byte b) throws IOException;
public synchronized void packBytes (byte[] b) throws IOException;
public synchronized void packBoolean (boolean b) throws IOException;
public synchronized void packChar (char c) throws IOException;
public synchronized void packShort (short s) throws IOException;
public synchronized void packInt (int i) throws IOException;
public synchronized void packLong (long l) throws IOException;
public synchronized void packFloat (float f) throws IOException;
public synchronized void packDouble (double d) throws IOException;
public synchronized void packString (String s) throws IOException;
};

public class InputBuffer
{
public InputBuffer ();

public final synchronized boolean valid ();
public synchronized byte[] buffer();
public synchronized int length ();

    /* unpack operations for standard Java types */

public synchronized byte unpackByte () throws IOException;
public synchronized byte[] unpackBytes () throws IOException;
public synchronized boolean unpackBoolean () throws IOException;
public synchronized char unpackChar () throws IOException;
public synchronized short unpackShort () throws IOException;
public synchronized int unpackInt () throws IOException;
public synchronized long unpackLong () throws IOException;
public synchronized float unpackFloat () throws IOException;
public synchronized double unpackDouble () throws IOException;
public synchronized String unpackString () throws IOException;
};
```

Uid

```
public class Uid implements Cloneable
{
public Uid ();
```

```

public Uid (Uid copyFrom);
public Uid (String uidString);
public Uid (String uidString, boolean errorsOk);
public synchronized void pack (OutputBuffer packInto) throws IOException;
public synchronized void unpack (InputBuffer unpackFrom) throws
IOException;

public void print (PrintStream strm);
public String toString ();
public Object clone () throws CloneNotSupportedException;
public synchronized void copy (Uid toCopy) throws UidException;
public boolean equals (Uid u);
public boolean notEquals (Uid u);
public boolean lessThan (Uid u);
public boolean greaterThan (Uid u);

public synchronized final boolean valid ();
public static synchronized Uid nullUid ();
};

```

AtomicAction

```

public class AtomicAction
{
public AtomicAction ();

public void begin () throws SystemException, SubtransactionsUnavailable,
NoTransaction;
public void commit (boolean report_heuristics) throws SystemException,
NoTransaction, HeuristicMixed,
HeuristicHazard, TransactionRolledBack;
public void rollback () throws SystemException, NoTransaction;
public Control control () throws SystemException, NoTransaction;
public Status get_status () throws SystemException;
/* Allow action commit to be suppressed */
public void rollbackOnly () throws SystemException, NoTransaction;

public void registerResource (Resource r) throws SystemException, Inactive;
public void registerSubtransactionAwareResource
(SubtransactionAwareResource sr)
throws SystemException,
NotSubtransaction;
public void registerSynchronization (Synchronization s) throws
SystemException,
Inactive;
};

```

Index

ArjunaTS		Nested top-level transactions	36
advanced programming	9	Nested transactions	35
class hierarchy	15	Object identity.....	10, 23
object models.....	23	Object serialisation.....	41
Asynchronous commit	36	Object state representation.....	10
Asynchronous prepare	36	Object storage	10
AtomicTransaction.....	14	Object store	
Checked transactions	33	further information	20
CheckedAction class.....	33	overview	20
Concurrency control policy.....	30	selecting.....	21
Configurable options.....	61	Object store types	
Core Classes		default implementation.....	65
<i>Buffer</i>	71	HashedStore.....	66
<i>Lock</i>	31	ShadowingStore.....	65
<i>LockManager</i>	29	ShadowNoFileLockStore.....	66
<i>ObjectState</i>	19, 70, 71	Object types	21
<i>ObjectStore</i>	20	OutputBuffer	18
<i>StateManager</i>	21	overview	18
Crash recovery		OutputObjectState.....	19
save_state and restore_state. 12, 22, 25, 41		overview	19
Creating objects	31	Persistent object lifecycle	12
Destroying objects	31	Persistent state.....	10
Identifying objects	10	issues.....	41
InputBuffer.....	18	Property variables	
overview	18	ASYNC_COMMIT	36
InputObjectState	19	ASYNC_PREPARE.....	36
overview	19	ENABLE_STATISTICS	34
Lifecycle of a persistent object	12	HASHED_DIRECTORIES.....	66
Lock store		JDBC2_USER_DB_ACCESS.....	67
further information	27	LOCKSTORE_DIR.....	28
implementations.....	13	LOCKSTORE_TYPE.....	28
overview	27	OBJECTSTORE_SYNC.....	20
selecting.....	28	OBJECTSTORE_TYPE.....	21, 63
LockManager	13, 28	OTS_TX_REAPER_MODE	39
Lock Conflicts	30	OTS_TX_REAPER_TIMEOUT	39
locking policy	30	releaselock.....	30
releaselock	30	restore_state	22, 25, 41
setlock.....	28, 30	Example.....	27
Examples	30	super.restore_state	12, 22, 25, 41

save_state	11, 22, 25, 41	Transaction context propagation.....	54
example.....	11, 26	Transactional Objects for Java.....	9
super.save_state	12, 22, 25, 41	class hierarchy	9
saving and restoring object states	18	configuration.....	61
setlock	16, 28	example.....	14, 37
State Management		Transactions	
deleting persistent objects.....	38	changing default timeout value.....	39
StateManager	11, 21	independent top level transactions.....	36
destroy	23	nested top-level transactions.....	36
direct use.....	42	nested transactions.....	35
<i>restore_state</i>	25	object constructors	40
<i>save_state</i>	25	timeouts	38
<i>terminate</i>	32	Use within <i>save_state</i> and <i>restore_state</i>	37
<i>type</i>	26	<i>type</i>	26
Subtransactions	35	<i>Example</i>	27
Threading		Type specific concurrency control.....	13
optimisations.....	36	Uid.....	10