



IMu Documentation

Using KE IMu's Java API

Document Version 1.2

EMu Version 4.0
IMu Version 1.0.03



Contents

SECTION 1	Introduction	1
SECTION 2	Using IMu's Java library	3
	Test Program	4
	Exceptions	5
SECTION 3	Connecting to an IMu server	7
	Handlers	8
SECTION 4	Accessing an EMu Module	9
	Searching a Module	10
	findKey	11
	findKeys	12
	findTerms	13
	Examples	14
	findWhere	16
	Number of matches	16
	Getting Information from Matching Records	17
	flag and offset	18
	count	19
	columns	20
	Return Values	21
	Attachments	24
	Reverse Attachments	25
	Rename a Column	26
	Grouping a set of nested table columns	27
	Column Sets	29
	A Simple Example	30
	Sorting	33
	keys	33
	flags	34
	Return Value	37
	Example	38
SECTION 5	Multimedia	41
SECTION 6	Maintaining State	47
	Example	50
SECTION 7	Exceptions	55
SECTION 8	Reference	57
	Class Handler	57
	Constructors	57
	Properties	58
	Methods	59

Class IMu	60
Class constants	60
Class IMuException	61
Constructors	61
Properties	61
Methods	62
Class Map	63
Methods	63
Class Module	66
Constructors	66
Properties	66
Methods	67
Class ModuleFetchResult	73
Properties	73
Class ModuleSortResult	74
Properties	74
Class ModuleSortTerm	75
Properties	75
Class Session	76
Class Properties	76
Constructors	76
Properties	77
Methods	78
Class TempInputStream	79
Constructors	79
Methods	79
Class Terms	80
Constructors	80
Properties	80
Methods	81
Enum TermsKind	82
Members	82
Index	83

SECTION 1

Introduction

IMu, or Internet Museum, broadly describes KE Software's strategy and toolset for distributing data held within EMu via the Internet. Distribution includes the publishing of content on the web, but goes far beyond this to cover sharing of data via the Internet (portals, online partnerships, etc.); publishing content to new mobile technologies; iPod guided tours, etc.

To facilitate these various Internet projects, KE has produced a set of documents that describe how to implement and customise IMu components, including:

- APIs (for Developers)
- Web pages for publishing EMu
- Tools, including:
 - iPhone / mobile interfaces
 - iPod guided tours

This document describes use of the IMu Java API.

SECTION 2

Using IMu's Java library

A single jar file, `imu-1-0-03.jar` (or higher), is required to develop an IMu-based application. This jar contains all the classes that make up the IMu Java API.

As with all jar files, the IMu jar must be included in the Java class path so that the java compiler and java runtime environment can find and use the IMu classes. Tools for Java development such as Eclipse and NetBeans allow you to add a reference to the IMu jar to your project.

All classes in the IMu Java API are included in the one package, `com.kesoftware.imu`. As is usual in Java development you are able to refer to an IMu class in your code by:

- Using the fully qualified name:

```
com.kesoftware.imu.Session session = new  
    com.kesoftware.imu.Session();
```

- Importing the required class explicitly:

```
import com.kesoftware.imu.Session;  
...  
Session session = new Session();
```

- Importing all the classes from the package implicitly:

```
import com.kesoftware.imu.*;  
...  
Session session = new Session();
```

Test Program

Compiling and running this very simple IMu program is a good test of whether the development environment has been set up properly for using IMu:

```
import com.kesoftware.imu.*;

class Hello
{
    static void main(String[] args)
    {
        System.out.format("IMu version %s%n", IMu.VERSION);
    }
}
```

The IMu library includes a class called `IMu`. This class includes a static `String` member called `VERSION` which contains the version of this IMu release.

Exceptions

Many of the methods in the IMu library objects throw exceptions when an error occurs. For this reason, code that uses IMu library objects should be surrounded with a `try/catch` block.

The following code is a basic template for writing Java programs which uses the IMu library:

```
import com.kesoftware.imu.*;
...
try
{
    // Create and use IMu objects
    ...
}
catch (Exception e)
{
    // Handle or report error
    ...
}
```

Most IMu exceptions throw an `IMuException` object. `IMuException` is a subclass of the standard Java `Exception`. In many cases your code can simply catch the standard `Exception` (as in this template). If more information is required about the exact `IMuException` thrown, see *Exceptions* (page 55).



Many of the examples that follow assume that code fragments have been surrounded with code structured in this way.

SECTION 3

Connecting to an IMu server

Most IMu based programs begin by creating a connection to an IMu server. Connections to a server are created and managed using IMu's `Session` class. Before connecting, both the name of the host and the port number to connect on must be specified. This can be done in one of three ways.

The simplest way to create a connection to an IMu server is to pass the `host` name and `port` number to the `Session` constructor and then call the `connect` method. For example:

```
import com.kesoftware.imu.Session;
...
Session mySession = new Session("server.com", 12345);
mySession.connect();
```

Alternatively, pass no values to the constructor and then set the `host` and `port` properties (using `setHost` and `setPort`) before calling `connect`:

```
import com.kesoftware.imu.Session;
...
mySession = new Session();
mySession.setHost("server.com");
mySession.setPort(12345);
mySession.connect();
```

If either the `host` or `port` is not set, the `Session` class default value will be used. These defaults can be overridden by setting the class (static) properties `defaultHost` and `defaultPort`:

```
import com.kesoftware.imu.Session;
...
Session.setDefaultHost("server.com");
Session.setDefaultPort(12345);
mySession = new Session();
mySession.connect();
```

This technique is useful when planning to create several connections to the same server or when wanting to get a handler object (page 8) to create the connection automatically.

Handlers

Once a connection to an IMu server has been established, it is possible to create handler objects to submit requests to the server and receive responses.



When a handler object is created, a corresponding object is created by the IMu server to service the handler's requests.

All handlers are subclasses of IMu's `Handler` class.



You do not typically create a `Handler` object directly but instead use a subclass.

In this document we examine the most frequently used handler, `Module`, which allows you to find and retrieve records from a single EMu module.

SECTION 4

Accessing an EMu Module

A program accesses an EMu module (or table, the terms are used interchangeably) using a `Module` class. The name of the table to be accessed is passed to the `Module` constructor. For example:

```
import com.kesoftware.imu.Module;
...
Module parties = new Module("eparties", mySession);
```

This code assumes that a `Session` object called `mySession` has already been created. If a `Session` object is not passed to the `Module` constructor, a session will be created automatically using the `defaultHost` and `defaultPort` class properties. See *Connecting to an IMu Server* (page 7) for details.

Once a `Module` object has been created, it can be used to search the specified module and retrieve records.

Searching a Module

One of the following methods can be used to search for records within a module:

- `findKey`
- `findKeys`
- `findTerms`
- `findWhere`

findKey

The `findKey` method searches for a single record by its key. The key is of type `long`.

For example, the following code searches for a record with a key of 42 in the Parties module:

```
import com.kesoftware.imu.Module;
...
Module parties = new Module("eparties", mySession);
long hits = parties.findKey(42);
```

The method returns the number of matches found, which is either 1 if the record exists or 0 if it does not.

findKeys

The `findKeys` method searches for a set of key values. The keys are passed as an array of longs:

```
Module parties = new Module("eparties", mySession);  
long[] keys = { 52, 42, 17 };  
long hits = parties.findKeys(keys);
```

or as an `ArrayList`:

```
Module parties = new Module("eparties", mySession);  
ArrayList<Long> keys = new ArrayList<Long>();  
keys.add(52L);  
keys.add(42L);  
keys.add(17L);  
long hits = parties.findKeys(keys);
```

The method returns the number of records found.

findTerms

The `findTerms` method is the most flexible and powerful way to search for records within a module. It can be used to run simple single term queries or complex multi-term searches.

The terms are specified using a `Terms` object. Once a `Terms` object has been created, add specific terms to it (using the `add` method) and then pass the `Terms` object to the `findTerms` method. For example, to specify a Parties search for records which contain a *First Name* of John and a *Last Name* of Smith:

```
Terms search = new Terms();
search.add("NamFirst", "John");
search.add("NamLast", "Smith");
...
long hits = myModule.findTerms(terms);
```

There are several points to note:

1. The first argument passed to the `add` method element contains the name of the column or an alias in the module to be searched.
2. The second argument contains the value for which to search.
3. A comparison operator can be included as a third argument (see example 3 below).

The operator specifies how the value supplied as the second argument of the array should be matched. Operators are the same as those used in TexQL (see KE's TexQL documentation for details).

Specifying an operator is optional. If none is supplied, the operator defaults to `matches`. This is not a real TexQL operator, but is translated by the search engine as the most "natural" operator for the type of column being searched. For example, with text columns `matches` is translated as "contains" and with integer columns it is translated as "=".



Unless it is really necessary to specify an operator, consider using the `matches` operator, or better still supplying no operator at all as this allows the server to determine the best type of search.



The first element of each term may be the name of a search alias. A search alias associates a name with one or more actual columns. Aliases are created using the `addSearchAlias` or `addSearchAliases` methods.

Examples

1. To search for the name `Smith` in the *Last Name* field of the Parties module, the following term can be used:

```
Terms search = new Terms();
search.add("NamLast", "Smith");
```

2. Specifying search terms for other types of columns is straightforward. For example, to search for records inserted on April 4, 2011

```
Terms search = new Terms();
search.add("AdmDateInserted", "Apr 4 2011");
```

3. To search for records inserted before April 4, 2011, it is necessary to add an operator:

```
Terms search = new Terms();
search.add("AdmDateInserted", "Apr 4 2011", "<");
```

4. By default, the relationship between the terms is a Boolean AND. This means that to find records which match both a *First Name* containing `John` and a *Last Name* containing `Smith` the `Terms` object can be created as follows:

```
Terms search = new Terms();
search.add("NamFirst", "John");
search.add("NamLast", "Smith");
```

5. A `Terms` object where the relationship between the terms is a Boolean OR can be created by passing the enumeration value `TermsKind.OR` to the `Terms` constructor. This means that:

```
Terms search = new Terms(TermsKind.OR);
search.add("NamFirst", "John");
search.add("NamLast", "Smith");
```

specifies a search for records where either the *First Name* contains `John` or the *Last Name* contains `Smith`.

6. Combinations of AND and OR search terms can be created. The `addAnd` method adds a new set of AND terms to the original `Terms` object. Similarly the `addOr` method adds a new set of OR terms. To restrict the search for a *First Name* of `John` and a *Last Name* of `Smith` to matching records inserted before April 4, 2011 or on May 1, 2011, specify:

```
Terms search = new Terms();
search.add("NamFirst", "John");
search.add("NamLast", "Smith");
Terms dates = search.addOr();
dates.add("AdmDateInserted", "Apr 4 2011", "<");
dates.add("AdmDateInserted", "May 1 2011");
```

7. To run a search, pass the terms object to the `findTerms` method:

```
Module parties = new Module("eparties", mySession);
Terms search = new Terms();
search.add("NamLast", "Smith");
long hits = parties.findTerms(myTerms);
```

As with other `find` methods, the return value contains the estimated number of matches.

8. To use a search alias, call the `addSearchAlias` method to associate the alias with one or more real column names before calling `findTerms`. Suppose we want to allow a user to search the Catalogue module for keywords. Our definition of a keywords search is to search the `SummaryData`, `CatSubjects_tab` and `NotNotes` columns. We could do this by building an OR search:

```
String keyword = ...;
...
Terms search = new Terms(TermsKind.OR);
search.add("SummaryData", keyword);
search.add("CatSubjects_tab", keyword);
search.add("NotNotes", keyword);
```

Another way of doing this is to register the association between the name `keywords` and the three columns we are interested in and then pass the name `keywords` as the column to be searched:

```
String keyword = ...;
...
Module catalogue = new Module("ecatalogue", mySession);
String[] columns =
{
    "SummaryData",
    "CatSubjects_tab",
    "NotNotes"
};
catalogue.addSearchAlias("keywords", columns);
...
Terms search = new Terms();
search.add("keywords", keyword);
catalogue.findTerms(search);
```

An alternative to passing the columns as an array of strings is to pass a single string, with the column names separated by semi-colons:

```
String keyword = ...;
...
Module catalogue = new Module("ecatalogue", mySession);
String columns = "SummaryData;CatSubjects_tab;NotNotes";
catalogue.addSearchAlias("keywords", columns);
...
Terms search = new Terms();
search.add("keywords", keyword);
catalogue.findTerms(search);
```

The advantage of using a search alias is that once the alias is registered a simple name can be used to specify a more complex OR search.

9. To add more than one alias at a time, use the `IMu Map` class to build an associative array of names and columns and call the `addSearchAliases` method:

```
Map aliases = new Map();
aliases.put("keywords",
    "SummaryData;CatSubjects_tab;NotNotes");
aliases.put("title", "SummaryData;TitMainTitle");
catalogue.addSearchAliases(aliases);
```

findWhere

With the `findWhere` method it is possible to submit a complete TexQL where clause.

```
Module parties = new Module("eparties", mySession);
String where = "NamLast contains 'Smith'";
long hits = parties.findWhere(where);
```

Although this method provides complete control over exactly how a search is run, it is generally better to use `findTerms` to submit a search rather than building a where clause. There are (at least) two reasons to prefer `findTerms` over `findWhere`:

1. Building the `where` clause requires the code to have detailed knowledge of the data type and structure of each column. The `findTerms` method leaves this task to the server. For example, specifying the term to search for a particular value in a nested table is straightforward. To find Parties records where the *Roles* nested table contains Artist, `findTerms` simply requires:

```
myTerms.add("NamRoles_tab", "Artist")
```

On the other hand, the equivalent TexQL clause is:

```
exists(NamRoles_tab where NamRoles contains 'Artist')
```

The TexQL for double nested tables is even more complex.

2. More importantly, `findTerms` is more secure.

With `findTerms` a set of terms is submitted to the server which then builds the TexQL where clause. This makes it much easier for the server to check for terms which may contain SQL-injection style attacks and to avoid them.

If your code builds a where clause from user entered data so it can be run using `findWhere`, it is much more difficult, if not impossible, for the server to check and avoid SQL-injection. The responsibility for checking for SQL-injection becomes yours.

Number of matches

All the `find` methods return the number of matches found by the search. For `findKey` and `findKeys` this number is always the exact number of matches found. The number returned by `findTerms` and `findWhere` is best thought of as an estimate. This estimate is almost always correct but because of the nature of the indexing used by the server's data engine (Texpress) the number can sometimes be an over-estimate of the real number of matches. This is similar to the estimated number of hits returned by a Google search.

Getting Information from Matching Records

Module's `fetch` method is used to get information from the matching records once the search of a module has been run. The server maintains the set of matching records in a list and `fetch` can be used to retrieve any information from any contiguous block of records in the list.

The simplest form of the `fetch` method takes four arguments:

- `flag`
- `offset`
- `count`
- `columns`



There are many different versions of the `fetch` method. See *Reference* (page 57) for details of each one.

flag and offset

The `flag` and `offset` arguments define the starting position of the block records to be fetched. The `flag` argument is a `String` and must be one of:

- `"start"`
- `"current"`
- `"end"`

The `"start"` and `"end"` flags refer to the first record and the last record in the matching set. The `"current"` flag refers to the position of the last record fetched by the previous call to `fetch`. If `fetch` has not been called, `"current"` refers to the first record in the matching set.

The `offset` argument is an `int`. It adjusts the starting position relative to the `flag`. A positive value for `offset` specifies a start after the position specified by `flag` and a negative value specifies a start before the position specified by `flag`.

For example, calling `fetch` with a `flag` of `"start"` and `offset` of 3 will cause `fetch` to return records starting from the fourth record in the matching set. Specifying a `flag` of `"end"` and an `offset` of -8 will cause `fetch` to return records starting from the ninth last record in the matching set.

To retrieve the next record after the last returned by the previous `fetch`, you would pass a `flag` of `"current"` and an `offset` of 1.

count

The `count` argument specifies the maximum number of records to be retrieved.

Passing a `count` value of 0 is valid. This causes `fetch` to change the current record without actually retrieving any data.

Using a negative value of `count` is also valid. This causes `fetch` to return all the records in the matching set from the starting position (specified by `flag` and `offset`).

columns

The `columns` argument is used to specify which columns should be included in the returned records. The argument can be either a simple `String`, an array of `Strings` or an `ArrayList` of `Strings`. In its simplest form each `String` contains a single column name, or several column names separated by semi-colons or commas.

For example, to retrieve the information for both the *NamFirst* and *NamLast* columns, you would do one of:

```
Module parties = new Module("eparties", mySession);
String columns = "NamFirst;NamLast";
parties.fetch("start", 0, 1, columns);
```

-OR-

```
String[] columns =
{
    "NamFirst",
    "NamLast"
};
parties.fetch("start", 0, 1, columns);
```

-OR-

```
import java.util.ArrayList;
...
ArrayList<String> columns = new ArrayList<String>();
columns.add("NamFirst");
columns.add("NamLast");
parties.fetch("start", 0, 1, columns);
```


Return Values

The `fetch` method returns records requested in a `ModuleFetchResult` object. This object contains three read-only properties:

- `count` (an `int`, accessed using `getCount()`)
- `hits` (a `long`, accessed using `getHits()`)
- `rows` (a `Map[]`, accessed using `getRows()`)

The `count` property is the number of records returned by the `fetch` request.

The `hits` property is the estimated number of matches in the result set. This number is returned for each `fetch` because the estimate can decrease as records in the result set are processed by the `fetch` method.

The `rows` property is an array containing the set of records requested. Each element of the `rows` array is itself a `Map` object. Each `Map` object contains entries for each column requested.

The `IMuMap` class is a subclass of Java's standard `HashMap`. It defines its key type to be `String`. It also provides some convenience methods for converting the types of elements stored in the map. See *Reference (page 57)* for details.

The following example shows a simple search of the EMu Parties module using `findTerms` with `fetch` used to retrieve a set of records:

```
import com.kesoftware.imu.*;
...
try
{
    Session mySession = new Session("server.com", 12345);

    Module parties = new Module("eparties", mySession);

    // Find all party records where Last Name contains 'smith'
    Terms search = new Terms();
    search.add("NamLast", "Smith");
    long hits = parties.findTerms(search);

    // We want to fetch the irn, NamFirst and NamLast
    // columns for each record.
    String[] columns =
    {
        "irn",
        "NamFirst",
        "NamLast"
    };

    // Fetch the first three records (at most) from the start
    // of the result set.
    ModuleFetchResult result = parties.fetch("start", 0, 3,
        columns);
    System.out.format("count: %d\n", result.getCount());
    System.out.format("hits: %d\n", result.getHits());
    System.out.format("rows:\n");
    Map[] rows = result.getRows();
    for (int i = 0; i < rows.length; i++)
    {
        Map row = rows[i];
        int rownum = row.getInt("rownum");
        long irn = row.getLong("irn");
        String first = row.getString("NamFirst");
        String last = row.getString("NamLast");

        System.out.format(" [%d]\n", i);
        System.out.format(" rownum: %d\n", rownum);
        System.out.format(" irn: %d\n", irn);
        System.out.format("   NamFirst: %s\n", first);
        System.out.format("   NamLast: %s\n", last);
    }
}
catch (Exception e)
{
    ...
}
```

The output of this code will be similar to:

```
count: 3
hits: 12
rows:
  [0]
    rownum: 1
    irn: 722
    NamFirst: Chris
    NamLast: SMITH
  [1]
    rownum: 2
    irn: 723
    NamFirst: Brad
    NamLast: Smith
  [2]
    Rownum: 3
    irn: 724
    NamFirst: Sylvia
    NamLast: Smith
```

Notice that data for each row includes the `irn`, `NamFirst` and `NamLast` elements, which correspond to the columns requested. Also notice that a `rownum` element is included. This element contains the number of the record within the result set (starting from 1) and is always included in the retrieved records.

Nested tables are returned as arrays of `Strings`. For example, if a `columns` argument of:

```
"NamLast;NamFirst;NamRoles_tab"
```

is passed, the loop from the previous example can be modified as follows:

```
for (int i = 0; i < rows.length; i++)
{
    Map row = rows[i];
    ...
    String[] roles = row.getStrings("NamRoles_tab");
    for (int j = 0; j < roles.length; j++)
        System.out.format("    NamRoles_tab[%d]: %s\n", j, roles[j]);
}
```

The output of this code will be similar to:

```
rows:
  [0]
    rownum: 1
    irn: 722
    NamFirst: Chris
    NamLast: SMITH
    NamRoles_tab[0]: Lyricist
    NamRoles_tab[1]: Pianist
  ...
```

Attachments

The set of columns requested can be more than simple column names. Columns from modules which the current record attaches to can also be requested. For example, suppose that the Catalogue module documents the creator of an object as an attachment (to a record in the Parties module) in a column called `CatCreatorRef`. If the Catalogue module is searched, it is possible to get the creator's last name for each Catalogue record in the result set as follows:

```
"CatCreatorRef.NamLast"
```

This technique can be extended to get information for more than one column:

```
"CatCreatorRef.(NamTitle;NamLast;NamFirst)"
```

The values are returned in a nested Map:

```
for (int i = 0; i < rows.length; i++)
{
    Map row = rows[i];
    ...
    Map creator = row.getMap("CreCreatorRef");
    String first = creator.getString("NamFirst");
    String last = creator.getString("NamLast");

    System.out.format("  Creator First Name %s\n", first);
    System.out.format("  Creator Last Name %s\n", last);
}
```

Reverse Attachments

In addition to standard attachment columns, it is possible to request information from so-called reverse attachments. A reverse attachment refers to one or more records which attach to the current record.

For example, to retrieve information from a set of Catalogue records which attach to the current Parties record via the Catalogue's *CatCreatorRef* column, specify:

```
"<ecatalogue:CatCreatorRef>.(irn,TitMainTitle)"
```

The following code fragment retrieves Parties IRN 53 and displays the CatCreatorRef reverse attachments:

```
Module parties = new Module("eparties", session);
long hits = parties.findKey(53);

String[] columns =
{
    "irn",
    "NamFirst",
    "NamLast",
    "<ecatalogue:CatCreatorRef>.(irn,TitMainTitle)"
};

ModuleFetchResult result = parties.fetch("start", 0, 1, columns);
```

The reverse attachments are returned as an array of Maps:

```
Map[] rows = result.getRows();
for (int i = 0; i < rows.length; i++)
{
    Map row = rows[i];
    ...
    Map[] att = row.getMaps("ecatalogue:CatCreatorRef");
    for (int j = 0; j < att.length; j++)
    {
        System.out.format("Row %d, Reverse Attachment %d\n", i, j);

        String title = att[j].getString("TitMainTitle");

        System.out.format("    title: %s\n", title);
    }
}
```

Rename a Column

It is possible to rename any column when it is returned by adding the new name in front of the real column being requested, followed by an equals sign.

For example, to request data from the `NamLast` column but rename it as `last_name`, specify:

```
"last_name=NamLast "
```

The returned `Map` will contain an element called `last_name` rather than `NamLast`.

This is particularly useful for complicated reverse attachment names:

```
"objects=<ecatalogue:CatCreatorRef>.(SummaryData) "
```

Grouping a set of nested table columns

A set of nested table columns can be grouped. Grouping allows the association between the columns to be reflected in the structure of the data returned. Consider the `Contributors` grid on the Details tab of the Narratives module, which contains two columns:

- `NarContributorRef_tab`
which contains a set of attachments to records in the Parties module.
- `NarContributorRole_tab`
which contains the roles for the corresponding contributors.

Each column can be retrieved separately as follows:

```
Module narratives = new Module("enarratives", mySession);

narratives.findKey(2);

String[] columns =
{
    "irn",
    "NarTitle",
    "NarContributorRef_tab.SummaryData",
    "NarContributorRole_tab"
};

ModuleFetchResult result = narratives.fetch("start", 0, 1,
    columns);
Map[] rows = result.getRows();
for (int i = 0; i < rows.length; i++)
{
    Map row = rows[i];

    Maps[] names = row.getMaps("NarContributorRef_tab");
    for (int j = 0; j < names.length; j++)
    {
        String summary = names[j].getString("SummaryData");
        System.out.format("Name %d: %s\n", j, summary);
    }

    String[] roles = row.getStrings("NarContributorRole_tab");
    for (int j = 0; j < roles.length; j++)
        System.out.format("Role %d: %s\n", j, roles[j]);
}
```

This produces output such as:

```
Name 0: Rising, John
Name 1: Graham, Beverley
Role 0: Artist
Role 1: Author
```

Although this works fine, the relationship between the contributor and his or her role is unclear. Grouping can make the relationship far clearer.

To group the columns, surround them with square brackets:

```
"[NarContributorRef_tab.SummaryData,NarContributorRole_tab]"
```

With this single change the previous code fragment looks like this:

```
String[] columns =
{
    "irn",
    "NarTitle",
    "[NarContributorRef_tab.SummaryData,NarContributorRole_tab]"
};

ModuleFetchResult result = narratives.fetch("start", 0, 1,
    columns);
Map[] rows = result.getRows();
for (int i = 0; i < rows.length; i++)
{
    Map row = rows[i];

    Map[] group = row.getMaps("group1");
    for (int j = 0; j < group.length; j++)
    {
        Map contrib = group[j].getMap("NarContributorRef_tab");
        String name = contrib.getString("SummaryData");

        String role = group[j].getString("NarContributorRole_tab");

        System.out.format("Contributor %d: Name %s; Role %s\n", j,
            name, role);
    }
}
```

This produces output such as:

```
Contributor 0: Name Rising, John; Role Artist
Contributor 1: Graham, Beverley; Role Author
```

By default, the group is given a name of group1, group2 and so on, which can be changed easily enough:

```
"contributors=[NarContributorRef_tab.SummaryData,
    NarContributorRole_tab]"
```


Column Sets

Every time `fetch` is called and a set of columns to retrieve is passed, the IMu server must parse these columns and check them against the EMu schema. For complex column sets, particularly those involving several references or reverse references, this can take time.

If `fetch` will be called several times with the same set of columns, it is a good idea to register the set of columns once and then simply pass the name of the registered set each time `fetch` is called.

Module's `addFetchSet` method is used to register a set of columns. This method takes two arguments:

- The name of the column set.
- The set of columns to be associated with that name.

For example:

```
String[] columns =
{
    "irn",
    "NamFirst",
    "NamLast"
};
parties.addFetchSet("PersonDetails", columns);
```

This registers the set of columns with the IMu server and gives it the name `PersonDetails`. This name can then be passed to any call to `fetch` and the same set of columns will be returned:

```
parties.fetch("start", 0, 5, "PersonDetails");
```

More than one set can be registered at once using `addFetchSets`. Simply build an associative array containing each set:

```
Map sets = new Map();
sets.put("PersonDetails", "irn;NamFirst;NamLast");
sets.put("OrganisationDetails", "irn;NamOrganisation");
parties.addFetchSets(sets);
```

Using column sets is very useful when maintaining state (page 47).

A Simple Example

In this example we build a simple command-line based Java program to search the Parties module by *Last Name* and display the full set of results. The name to be searched for will be passed to the program as a command-line argument:

```

import com.kesoftware.imu.*;

public class Example
{
    public static void
    main(String[] args)
    {
        if (args.length != 1)
        {
            System.err.format("Usage: example name%n");
            System.exit(1);
        }
        try
        {
            process(args[0]);
        }
        catch (IMuException e)
        {
            System.err.format("Sorry, an error occurred: %s%n", e);
            System.exit(1);
        }
    }

    private static void
    process(String lastName) throws IMuException
    {
        Session mySession = new Session("server.com", 40999);
        mySession.connect();
        Module parties = new Module("eparties", mySession);

        // Build search term and run search
        Terms search = new Terms();
        search.add("NamLast", lastName);
        long hits = parties.findTerms(search);

        // Build list of columns to fetch
        String[] columns =
        {
            "NamFirst",
            "NamLast"
        };

        // Fetch all the matches in one go by passing count < 0
        ModuleFetchResult result = parties.fetch("start", 0, -1,
            columns);

        // Display the results
        System.out.format("Number of matches: %d%n",
            result.getHits());
        Map[] rows = result.getRows();
        for (int i = 0; i < rows.length; i++)
        {
            long rownum = rows[i].getLong("rownum");
            String first = rows[i].getString("NamFirst");
            String last = rows[i].getString("NamLast");
            System.out.format("%d: %s %s%n", rownum, first, last);
        }
    }
}

```

The results generated look like this:

```
Number of matches: 5
```

```
1: Percy JONES
```

```
2: Marilyn JONES
```

```
3: Lee Jones
```

```
4: David Jones
```

```
5: William Jones
```

Sorting

The matching set of results can be sorted using `Module's sort` method. This method takes two arguments:

- `keys`
- `flags`

keys

The `columns` argument is used to specify the columns by which to sort the result set. The argument can be either a simple `String`, an array of `Strings` or an `ArrayList` of `Strings`. Each string can be a simple column name or a set of column names, separated by semi-colons or commas. Each column name can be preceded by a `+` or `-`. A leading `+` indicates that the records should be sorted in ascending order. A leading `-` indicates that the records should be sorted in descending order.

For example, to sort a set of `Parties` records first by *Party Type* (ascending), then *Last Name* (descending) and then *First Name* (ascending):

```
String keys = "+NamPartyType;-NamLast;+NamFirst";
```

-OR-

```
String keys[] =
{
    "+NamPartyType",
    "-NamLast",
    "+NamFirst"
};
```

-OR-

```
ArrayList<String> keys = new ArrayList<String>();
keys.add("+NamPartyType");
keys.add("-NamLast");
keys.add("+NamFirst");
```



If a sort order (`+` or `-`) is not given, the sort order defaults to ascending.

flags

The `flags` argument is used to pass one or more flags to control the way the sort is carried out. As with the `keys` argument, the `flags` argument can be a simple `String`, an array of `Strings` or an `ArrayList` of `Strings`. Each string can be a single flag or a set of flags separated by semi-colons or commas.

The following flags control the type of comparisons used when sorting:

"word-based"	<p><code>sort</code> disregards all punctuation and white spaces (more than the one space between words). For example:</p> <p>Traveler's Inn</p> <p>will be sorted as</p> <p>Travelers Inn</p>
"full-text"	<p><code>sort</code> includes all punctuation and white spaces. For example:</p> <p>Traveler's Inn</p> <p>will be sorted as</p> <p>Traveler's Inn</p> <p>and will therefore differ from:</p> <p>Traveler's Inn</p>
"compress-spaces"	<p><code>sort</code> includes punctuation but disregards all white space (with the exception of a single space between words). For example:</p> <p>Traveler's Inn</p> <p>will be sorted as</p> <p>Traveler's Inn</p>



If none of these flags is included, the comparison defaults to "word-based".

The following flags modify the sorting behaviour:

"case-sensitive"	<p>sort is sensitive to upper and lower case. For example:</p> <p>Melbourne gallery</p> <p>will be sorted separately to</p> <p>Melbourne Gallery</p>
"order-insensitive"	<p>Values in a multi-value field will be sorted alphabetically regardless of the order in which they display. For example, a record which has the following values in the <i>NamRoles_tab</i> column in this order:</p> <p>Collection Manager</p> <p>Curator</p> <p>Internet Administrator</p> <p>and another record which has the values in this order:</p> <p>Internet Administrator</p> <p>Collection Manager</p> <p>Curator</p> <p>will be sorted the same.</p>
"null-low"	<p>Records with empty records will be placed at the start of the result set rather than at the end.</p>
"extended-sort"	<p>Values that include diacritics will be sorted separately to those that do not. For example, <i>entrée</i> will be sorted separately to <i>entree</i>.</p>

The following flags can be used when generating a summary of the sorted records:

"report"	<p>A summary of the sort is generated. The summary is contained in a <code>ModuleSortResult</code> object. The result is hierarchically structured, summarising the number of records which match each of the sort keys. See the example (page 38) for an illustration of the structure.</p>
"table-as-text"	<p>All data from multi-valued columns will be treated as a single value (joined by line break characters) in the summary results array.</p> <p>For example, for a record which has the following values in the <i>NamRoles_tab</i> column:</p> <p>Collection Manager, Curator, Internet Administrator</p> <p>the summary will include statistics for a single value:</p> <p>Collection Manager Curator Internet Administrator</p> <p>Thus the number of values in the summary results display will match the number of records.</p> <p>If this option is not included, each value in a multi-valued column will be treated as a distinct value in the summary. Thus there may be many more values in the summary results than there are records.</p>

Return Value

The `sort` method returns null unless the `report` flag is used.

If the `report` flag is used, the `sort` method returns a `ModuleSortResult` object. This object contains two read-only properties:

- `count` (an `int`, accessed using `getCount()`).
- `terms` (an array of `ModuleSortTerm` objects, accessed using `getTerms()`).

The `count` property is the number of distinct terms in the primary sort key.

The `terms` property is an array containing the list of distinct terms associated with the primary key in the sorted result set.

Each element in the `terms` array is a `ModuleSortTerm` object. This object contains three read-only properties which describe the term:

- `value` (a `String`, accessed using `getValue()`).
- `count` (a `long`, accessed using `getCount()`).
- `nested` (a `ModuleSortResult` object, accessed using `getNested()`).

The `value` property is the distinct value itself.

The `count` property is the number of records in the result set which have this value.

The `nested` property is a nested `ModuleSortResult` object. This holds values for secondary sorts within the primary sort. This is illustrated in the following example:

Example

In this example we run a three-level sort on a set of Parties records, sorting first by *Party Type*, then *Last Name* (descending) and then by *First Name*. Setting up and running the sort is straightforward:

```
Module parties = new Module("eparties", ...);  
...  
parties.findTerms(...);  
...  
String[] keys =  
{  
    "+NamPartyType",  
    "-NamLast",  
    "+NamFirst"  
};  
String[] flags =  
{  
    "full-text",  
    "case-sensitive",  
    "report"  
};  
ModuleSortResult result = parties.sort(keys, flags);
```

We can write a simple method to display the result summary. This example displays the distinct terms (and their counts) for the primary sort key (*Party Type*). Nested within each primary term is the set of distinct terms for the secondary key (*Last Name*) and nested within this list is the set of distinct terms for the tertiary key (*First Name*).

This is most simply done by making the display method recursive. The `showSummary` method below illustrates how to walk the `ModuleSortResult` structure:

```

private void
showSummary(ModuleSortResult result, int indent)
{
    // Build a prefix string to indent the data correctly
    String prefix = "";
    for (int i = 0; i < indent; i++)
        prefix += "  ";

    // Display each term at this level
    ModuleSortTerm[] terms = result.getTerms();
    for (int i = 0; i < terms.length; i++)
    {
        ModuleSortTerm term = terms[i];

        // Get the value and count properties for the term
        String value = term.getValue();
        int count = term.getCount();

        // Print them out, indented appropriately
        System.out.format("%s%2d. \"%s\" (%d)%n",
            prefix, i, value, count);

        // If the nested property is defined then there are
        // further values for secondary, tertiary keys and so on
        // so we call showSummary recursively.
        ModuleSortResult nested = term.getNested();
        if (nested != null)
            showSummary(nested, indent + 1);
    }
}

```

This will produce output similar to the following:

```

"Person" (2086)
  0. "Young" (4)
    0. "Derek" (1)
    1. "Don" (1)
    ...
  1. "Williams" (5)
    0. "Arthur" (1)
    1. "John" (2)
    ...

```


SECTION 5

Multimedia

The multimedia resources associated with an EMu record can be retrieved using Module's `fetch` method by specifying a special column called `multimedia`. When this column is requested the server returns the set of multimedia attachments associated with the record in question.

The set is returned as an array of `Map` objects. Each map includes the following information:

- `irn`
The `irn` of the resource in EMu's Multimedia module.
- `type`
The media type: typically `image`, `audio`, `video`, etc.
- `format`
The media format or sub-type such as `jpeg` or `tiff` for image formats, `wav` or `mpeg` for audio.

This is equivalent to the column request:

```
multimedia=MulMultiMediaRef_tab.
(
  irn,
  type=MulMimeType,
  format=MulMimeFormat
)
```

with the addition that the result does not contain any empty entries (i.e. entries corresponding to null values in the *MulMultiMediaRef_tab* column) or any entries for Multimedia records which are not accessible via IMu.

For example:

```
Session mySession = new Session("server.com", 40999);
mySession.connect();

Module parties = new Module("eparties", mySession);

// Build the search and run it
Terms search = new Terms();
search.add("NamLast", "Pavarotti");
parties.findTerms(search);

// Build list of columns to fetch
String[] columns =
{
  "NamFirst",
  "NamLast",
  "multimedia"
};
```

```
// We are only interested in the first record
ModuleFetchResult result = parties.fetch("start", 0, 1, columns);
Map[] rows = result.getRows();
Map row = rows[0];

// Display the results
String first = row.getString("NamFirst");
String last = row.getString("NamLast");
Map[] multimedia = row.getMaps("multimedia");

System.out.format("First Name: %s\n", first);
System.out.format("Last Name: %s\n", last);
System.out.format("multimedia (%d)\n", multimedia.length);
for (int i = 0; i < multimedia.length; i++)
{
    Map entry = multimedia[i];

    long irn = entry.getLong("irn");
    String type = entry.getString("type");
    String format = entry.getString("format");

    System.out.format("    irn %d: %s/%s\n", irn, type, format);
}
```

will produce out such as:

```
First Name: Luciano
Last Name: PAVAROTTI
multimedia (11)
    irn 100096: image/gif
    irn 100100: image/gif
    irn 100101: image/gif
    irn 100102: image/gif
    irn 100105: image/jpeg
    irn 100095: video/quicktime
    irn 100103: video/quicktime
    irn 100098: audio/wav
    irn 100099: audio/wav
    irn 100104: audio/wav
    irn 100097: application/msword
```

The `multimedia` column is an example of an IMu "virtual" column. The column does not actually exist in the EMu table being accessed. Instead, the IMu server interprets the request for the column and builds an appropriate response. There are other virtual columns that can be used when accessing a record's multimedia attachments:

- `images`
This returns the subset of multimedia attachments which have a mime type of `image`. Like `multimedia`, this is returned as an array of `Map` objects.
- `image`
The preferred image from the set of images. Currently this is the same as the first entry in the array returned by `images`. However, future versions of EMu may allow another multimedia attachment to be flagged as the preferred image, in which case the `image` column will return information for the preferred resource, rather than the first one. This is returned as a single `Map` object.

- `videos`
This returns the subset of multimedia attachments which have a mime type of `video`.
- `video`
The preferred video from the set of `videos`. Currently this is the same as the first entry in the array returned by `videos`.

All these virtual columns act as reference columns into the Multimedia module. This means that other Multimedia columns can also be requested from the corresponding Multimedia record. For example, to include the publisher (*DetPublisher*) in the information returned for each attached multimedia resource:

```
multimedia.DetPublisher
```

The returned `Maps` will include a `DetPublisher` entry as well as the standard `irn`, `type` and `format` entries.

Any standard columns from the Multimedia module can be requested in this way. In addition, the Multimedia module includes a virtual column, `resource`, which can be used get access to the contents of the actual multimedia resource. The `resource` column is returned as a `Map` object. The object includes the following information:

- `identifier`
The contents of the multimedia record's *MulIdentifier* field.
- `contentType`
The media type: typically `image`, `audio`, `video`, etc.
- `contentType`
The media format or sub-type such as `jpeg` or `tiff` for image formats, `wav` or `mpeg` for audio.
- `size`
The size of the resource in bytes.
- `file`
A `TempInputStream` object. This provides a read-only handle to a temporary copy of the resource itself. The `TempInputStream` class is an IMu-specific subclass of Java's standard `FileInputStream` and the standard input methods can be used to read the contents of the resource. The temporary copy of the file is discarded when the stream is closed or destroyed.
- `height`
For images, the height of the image in pixels.
- `width`
For images, the width of the image in pixels.

The following code fragment retrieves Parties IRN 53, displays the information for its preferred attached image and creates a copy of the resource in a file called `image-copy`:

```
Module parties = new Module("eparties", mySession);
long hits = parties.findKey(53);

String[] columns =
{
    "NamFirst",
    "NamLast",
    "image.resource"
};

ModuleFetchResult result = parties.fetch("start", 0, 1, columns);
...
Map[] rows = result.getRows();

// Because we did a findKey() search, we are only
// interested in the first row.
Map row = rows[0];

Map image = row.getMap("image");
Map resource = image.getMap("resource");

// Print out information about the resource
String identifier = resource.getString("identifier");
String mimeType = resource.getString("mimeType");
String mimeTypeFormat = resource.getString("mimeTypeFormat");
long size = resource.getLong("size");

System.out.format("identifier: %s\n", identifier);
System.out.format("mimeType: %s\n", mimeType);
System.out.format("mimeTypeFormat: %s\n", mimeTypeFormat);
System.out.format("size: %d\n", size);

// Save a copy of the resource
FileInputStream temp = (FileInputStream) resource.get("file");
FileOutputStream copy = new FileOutputStream("image-copy");
byte[] buffer = new byte[4096]; // 4K buffer
while (temp.read(buffer) > 0)
    copy.write(buffer);
copy.close();
```

This will produce output similar to:

```
identifier: LucianoPavarotti.gif
mimeType: image
mimeTypeFormat: gif
size: 19931
```

as well as creating a file called `image-copy` which contains the copy of the image itself.

The previous example retrieves a binary copy of the master resource in its original format. It is also possible to modify how the resource is returned. This is done by adding modifiers to the `resource` column request. Modifiers are added after the column name and inside a set of braces.

The modifiers which can be applied to the `resource` column are:

- `encoding`

Specifies that the resource returned should be encoded. The only currently supported encoding is `base64`. By default the resource is returned as raw binary data.

Example:

```
resource{encoding:base64}
```

- `checksum`

Specifies that the information returned with the resource should include a checksum. The checksum requested can be `crc32` or `md5`.

Example:

```
resource{checksum:crc32}
```

In addition other modifiers can be applied to image resources:

- `format`

Specifies the format of the required image. If the master image is already in the required format, then it is returned. Otherwise the image is reformatted on-the-fly and the reformatted image is returned.

Example:

```
resource{format:gif}
```

This requests that the image is returned as a gif.

The IMu server uses ImageMagick to process the image and the range of supported formats is very large. The complete list is available from:

<http://www.imagemagick.org/script/formats.php>

- `height`

Specifies the height of the image required in pixels. If the record contains a resolution with this height, this resolution is returned. Otherwise the closest matching larger resolution is resized to the requested height on-the-fly and the resized image is returned.

Example:

```
resource{height:200}
```

- `width`

Specifies the width of the image required in pixels. If the record contains a resolution with this width, this resolution is returned. Otherwise the closest matching larger resolution is resized to the requested width on-the-fly and the resized image is returned.

Example:

```
resource{width:300}
```

- `bestfit`

If set to `yes`, the image returned is the existing resolution which most closely matches the specified height or width. No on-the-fly resizing is done.

Example:

```
resource{height:300,bestfit:yes}
```

This returns the image closest to, but larger than, 300 pixels high.

- `aspectratio`

Controls whether the image's aspect ratio should be maintained when both a height and a width are specified. If set to `no`, the aspect ratio is not maintained.

Example:

```
resource{height:300,width:300,aspectratio:no}
```

- `source`

Controls which image is used as the basis for any reformatting that is required.

By default, if no height or width is specified, the master is used as the source image. However, if a height or width is supplied, then by default the closest sized but larger resolution is used as the source. This saves processing time but may not produce the best result when dealing with lossy formats (such as jpeg). To override this, a source value of `master` can be specified.

Example:

```
resource{height:300,source:master}
```

This specifies that the image is generated by resizing the master to 300 pixels high, rather than by using any appropriate resolution.

The source value can also be `thumbnail`. In this case the image thumbnail is used as the source. Typically you would not want to apply size transformations to the thumbnail but this provides a simple way of retrieving the image's 90x90 thumbnail:

```
resource{source:thumbnail}
```

SECTION 6

Maintaining State

One of the biggest drawbacks of the earlier example (page 30) is that it fetches the full set of results at one time, which is impractical for large result sets. It is more practical to display a full set of results across multiple pages and allow the user to move forward or backward through the pages.

This is simple in a conventional application where a connection to the server is maintained until the user terminates the application. In a web implementation however, this seemingly simple requirement involves a considerably higher level of complexity due to the *stateless* nature of web pages. One such complexity is that each time a new page of results is displayed, the initial search for the records must be re-executed. This is inconvenient for the web programmer and potentially slow for the user.

The IMu server provides a solution to this. When a handler object is created, a corresponding object is created on the server to service the handler's request: this server-side object is allocated a unique identifier by the IMu server. When making a request for more information, the unique identifier can be used to connect a new handler to the same server-side object, with its state intact.

The following example illustrates the connection of a second, independently created `Module` object to the same server-side object:

```
// Create a module object as usual
Module first = new Module("eparties", session);

// Run a search - this will create a server-side object
long[] keys = { 1, 2, 3, 4, 5, 42 };
first.findKeys(keys);

// Get a set of results
ModuleFetchResult result1 = first.fetch("start", 0, 2,
    "SummaryData");

// Create a second module object
Module second = new Module("eparties", session);

// Attach it to the same server-side object as the
// first module. This is the key step.
second.setID(first.getID());

// Get a second set of results from the same search
ModuleFetchResult result2 = second.fetch("current", 1, 2,
    "SummaryData");
```

Although two completely separate `Module` objects have been created, they are each connected to the same server-side object by virtue of having the same `id` property. This means that the second `fetch` call will access the same result set as the first `fetch`. Notice that a flag of `current` has been passed to the second call.

The `current` state is maintained on the server-side object, so in this case the second call to `fetch` will return the third and fourth records in the result set.

While this example illustrates the use of the `id` property, it is not particularly realistic as it is unlikely that two distinct objects which refer to the same server-side object would be required in the same piece of code. The need to re-connect to the same server-side object when generating another page of results is far more likely. This situation involves creating a server-side `Module` object (to search the module and deliver the first set of results) in one request and then re-connecting to the same server-side object (to fetch a second set of results) in a second request. As before, this is achieved by assigning the same identifier to the `id` property of the object in the second page, but two other things need to be considered.

By default the IMu server destroys all server-side objects when a session finishes. This means that unless the server is explicitly instructed not to do so, the server-side object may be destroyed when the connection from the first page is closed. Telling the server to maintain the server-side object only requires that the `destroy` property on the object is set to `false` before calling any of its methods. In the example above, the server would be instructed not to destroy the object as follows:

```
Module module = new Module("eparties", session);
module.setDestroy(false);
long[] keys = { 1, 2, 3, 4, 5, 42 };
module.findKeys(keys);
```

The second point is quite subtle. When a connection is established to a server, it is necessary to specify the port to connect to. Depending on how the server has been configured, there may be more than one server process listening for connections on this port. Your program has no control over which of these processes will actually accept the connection and handle requests. Normally this makes no difference, but when trying to maintain state by re-connecting to a pre-existing server-side object, it is a problem.

For example, suppose there are three separate server processes listening for connections. When the first request is executed it connects, effectively at random, to the first process. This process responds to the request, creates a server-side object, searches the Parties module for the terms provided and returns the first set of results. The server is told not to destroy the object and passes the server-side identifier to another page which fetches the next set of results from the same search.

The problem comes when the next page connects to the server again. When the connection is established any one of the three server processes may accept the connection. However, only the first process is maintaining the relevant server-side object. If the second or third process accepts the connection, the object will not be found.

The solution to this problem is relatively straightforward. Before the first request closes the connection to its server, it must notify the server that subsequent requests need to connect explicitly to that process. This is achieved by setting the `Session` object's `suspend` property to `true` prior to submitting any request to the server:

```
Session session = new Session("server.com", 12345);  
Module module = new Module("eparties", session);  
...  
session.setSuspend(true);  
module.findKeys(...);
```

The server handles a request to suspend a connection by starting to listen for connections on a second port. Unlike the primary port, this port is guaranteed to be used only by that particular server process. This means that a subsequent page can reconnect to a server on this second port and be guaranteed of connecting to the same server process. This in turn means that any saved server-side object will be accessible via its identifier. After the request has returned (in this example it was a call to `findKeys`), the `Session` object's `port` property holds the port number to reconnect to:

```
session.setSuspend(true);  
module.findKeys(...);  
int reconnect = session.getPort();
```

Example

Although this may appear to be a little complicated, it is not in fact too difficult to manage in practice.

To illustrate we'll create a simple JSP-based website to display the list of matching names in blocks of five records per page. We'll provide simple **Next** and **Prev** links to allow the user to move through the results, and we will use some more GET parameters to pass the `port` we want to reconnect to, the identifier of the server-side object and the `rownum` of the first record to be displayed.

The main code is in `results.jsp`.

Rather than use a `try/catch` block, we specify an error page to catch exceptions (we will need a suitable error handling page - see below).

First we import the IMu libraries and set the error page directive:

```
<%@ page import="com.kesoftware.imu.*"
    errorPage="exceptionHandler.jsp" %>
```

Next we create the `IMuSession` object (do not call it `session` or it will clash with the existing JSP `HttpSession` object named `session`):

```
// create new session object
final Session imuSession = new Session();
imuSession.setHost(imuHost);
```

We set the `port` property to a standard value unless a `port` parameter has been passed in the URL:

```
if (request.getParameter("port") != null)
    imuPort = Integer.parseInt(request.getParameter("port"));
imuSession.setPort(imuPort);
```

Next we connect to the server. We immediately set the `suspend` property to `true` to tell the server that we may want to connect again (this ensures the server listens on a new, unique port):

```
// Establish connection and tell the server
** we may want to re-connect

imuSession.connect();
imuSession.setSuspend(true);
```

We then create the client-side `IMuModule` object and set its `destroy` property to `false`, ensuring the server will not destroy it:

```
final Module module = new Module("eparties", imuSession);
module.setDestroy(false);
```

If the URL includes a `name` parameter, we need to do a new search. Alternatively, if it includes an `id` parameter, we need to connect to an existing server-side object:

```
Long hits = 0L;
if (request.getParameter("name") != null)
{
    final Terms searchTerms = new Terms(TermsKind.OR);
    searchTerms.add("NamLast", request.getParameter("name"),
        "matches");
    hits = module.findTerms(searchTerms);
}

// Otherwise, if id is supplied reattach to
** existing server-side object

else if (request.getParameter("id") != null)
    module.setID(request.getParameter("id"));

// Otherwise, we can't process
else
    throw new Exception("no name or id");
```

As usual, we build a list of columns to fetch:

```
final String[] columns = { "NamFirst", "NamLast" };
```

If the URL includes a `rownum` parameter, fetch records starting from there. Otherwise start from record number 1:

```
// Work out which block of records to fetch
int rownum = 1;
if (request.getParameter("rownum") != null)
    rownum = Integer.parseInt(request.getParameter("rownum"));
```

Fetch the records:

```
// Fetch next five records
final ModuleFetchResult result =
    module.fetch("start", (rownum - 1), 5, columns);

// Save rows in convenient variable
** for later display

final Map[] rows = result.getRows();
final Integer port = imuSession.getPort();
final String id = module.getID();
```

Now we can build the main page.

```
<body>
  <!-- show hit count -->
  <p>Number of matches: <% out.print(hits); %></p>
  <table>
    <%
      // Display each match in a separate row in a table
      for (int i = 0; i < rows.length; i++)
      {
        final Map row = rows[i];
        Long rnum = row.getLong("rownum");
        out.println("<tr>");
        out.println("\t<td>" + rnum.toString() + "</td>");
        out.println("\t<td>" +
          row.getString("NamFirst") +
          " " + row.getString("NamFirst") +
          "</td>");
        out.println("</tr>");
      }
    %>
  </table>
  ...
```

Finally we add the **Prev** and **Next** links to allow the user to page backwards and forwards through the results. This is the most complicated part! First, we want to ensure that we connect to the same server and server-side object, so we add the appropriate port and id parameters to our URL:

```
<%
  // Add the Prev and Next links
  String url = request.getRequestURL().toString();
  url += "?port=" + port;
  url += "&id=" + id;
  ...
```

If we are not already showing the first record, we add a **Prev** link to allow the user to go back one page in the result set:

```
final Map first = rows[0];
final Long firstRowNum = first.getLong("rownum");

// If we are not already showing the first record,
** we add a Prev link

if (firstRowNum > 1)
{
  Long prev = (firstRowNum - 5);
  if (prev < 1L)
    prev = 1L;
  final String prevLink = url + "&rownum=" + prev;
  out.println("<a href=\"" +
    url + "&rownum=" + prev + "\">Prev</a>");
}
```

Similarly, if we are not already showing the last record, we add a **Next** link to allow the user to go forward one page:


```

final Map last = rows[rows.length - 1];
final Long lastRowNum = last.getLong("rownum");
// if we are not already showing the last record,
** we add a Next link

if (lastRowNum < hits)
{
    Long next = lastRowNum + 1L;
    final String nextLink = url + "&rownum=" + next;
    out.println("<a href=\"\" + url +
        "&rownum=" + next + "\">Next</a>");
}
%>
</body>

```

The resulting web page looks like this:



A simple error page might be:

```

<%@ page isErrorPage="true" import="java.io.*" %>
<html>
<head>
    <title>Error</title>
</head>
<body>
<h2>An Error occured...</h2>
<%= exception.toString() %><br>
<%
    out.println("<!--");
    StringWriter sw = new StringWriter();
    PrintWriter pw = new PrintWriter(sw);
    exception.printStackTrace(pw);
    out.print(sw);
    sw.close();
    pw.close();
    out.println("-->");
%>

</body>
</html>

```


SECTION 7

Exceptions

When an error occurs, the IMu Java API throws an exception. The exception is an `IMuException` object. This is a subclass of Java's standard `Exception` class.

For simple error handling all that is usually required is to catch the exception as an `Exception` object and report the exception as a string:

```
try
{
    ...
}
catch (Exception e)
{
    System.err.format("Error: %s%n", e);
    System.exit(1);
}
```

`IMuException` overrides the `Exception`'s `toString` and returns an error message.

To handle specific IMu errors it is necessary to catch the exception as an `IMuException` object. `IMuException` includes a property called `id`. This is a string and contains the internal IMu error code for the exception. For example, you may want to catch the exception raised when a `Session`'s `connect` method fails and try to connect to an alternative server:

```
String mainServer = "server1.com";
String alternativeServer = "server2.com";
Session mySession = new Session();
mySession.setPort = (...);

// Try the main server first
mySession.Host = mainServer;
try
{
    mySession.Connect();
}
catch (IMuException e)
{
    // Check for specific SessionConnect error
    if (e.getID() != "SessionConnect")
        throw;

    // Now try the alternative server
    mySession.setHost (alternativeServer);
    mySession.Connect();
}
// By the time we get to here the session is connected
// to either the main server or the alternative.
```

SECTION 8

Reference

Class Handler

`com.kesoftware.imu.Handler`

Provides a general low-level interface to creating server-side objects.

Constructors

```
public Handler(Session session)
```

Creates an object which can be used to interact with server-side objects.

Parameters

<code>session</code>	A <code>Session</code> object to be used to communicate with the IMu server.
----------------------	--

```
public Handler()
```

Same as `Handler` above but a new session is created automatically using the `Session` class's default host and port values.

Properties

Object create

getter `getCreate()`

setter `setCreate(Object create)`

An object to be passed to the server when the server-side object is created. To have any effect this must be set before any object methods are called. This property is usually only set by sub-classes of `Handler`.

boolean destroy

getter `getDestroy()`

setter `setDestroy(Boolean destroy)`

A flag controlling whether the corresponding server-side object should be destroyed when the session is terminated.

String id

getter `getID()`

setter `setID(String id)`

The unique identifier assigned to the server-side object once it has been created.

String language

getter `getLanguage()`

setter `setLanguage(String language)`

The language to be used in the server.

String name

getter `getName()`

setter `setName(String name)`

The name of the server-side object to be created. This must be set before any object methods are called.

Session session

getter `getSession()`

The session object used by the handler to communicate with the IMu server.

Methods

```
public Object call(String method, Object parameters)
```

Calls a method on the server-side object.

Parameters

method	The name of the method to be called.
parameters	Any parameters to be passed to the method. The <code>call</code> method uses Java's reflection to determine the structure of the parameters to be transmitted to the server.

Returns An object containing the result returned by the server-side method.

Throws `IMuException` if a server-side error occurred.

```
public Object call(String method)
```

Same as `call` above but without any additional parameters.

```
public Map request(Map request)
```

Submits a low-level request to the IMu server. This method is chiefly used by the `call` method above.

Parameters

request	A <code>Map</code> object containing the request parameters.
---------	--

Returns A `Map` object containing the server's response.

Throws `IMuException` if a server-side error occurred.

Class IMu

`com.kesoftware.imu.IMu`

Simple class containing general IMu properties. This class cannot be instantiated.

Class constants

`String` `VERSION`

The version number of the IMu API.

Class IMuException

`com.kesoftware.imu.IMuException`

Extends: `java.lang.Exception`

Class for IMu-specific exceptions.

Constructors

`public IMuException(String id, Object... args)`

Creates an IMu specific exception.

Parameters

<code>id</code>	A string exception code.
<code>args</code>	Any additional arguments used to provide further information about the exception.

`public IMuException(String id)`

Same as `IMuException` above but without any additional arguments.

Properties

`Object[]` `args`

getter `getArgs()`

setter `setArgs(Object[] args)`

A flag controlling whether the corresponding server-side object should be destroyed when the session is terminated.

`String` `id`

getter `getID()`

The unique identifier assigned to the server-side object once it has been created.

Methods

```
public String toString()
```

Overrides the standard `Object toString` method.

Returns A string description of the exception.

Class Map

`com.kesoftware.imu.Map`

Extends: `java.util.Arrays.HashMap<String, Object>`

Provides a simple map class with string keys and a set of convenience methods for getting values of certain types.

Methods

```
public <T> T[] getArray(String name, Class<T[]> type)
```

Gets the value associated with the key `name` and returns it as an array of the type specified by `type`.

Typical usage is:

```
long[] array = map.getArray("key", Long[].class);
```

Parameters

<code>name</code>	The key whose associated value is to be returned.
<code>type</code>	The type of the array required.

Returns The correctly typed array.

```
public boolean getBoolean(String name)
```

Gets the value associated with the key `name` and returns it as a `boolean`.

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns The value, interpreted as a `boolean`. Null values are considered false. Numeric values are considered false if they evaluate to zero and true otherwise. Any other non-boolean value is converted to a `String` and then parsed as a `Boolean`.

```
public double getDouble(String name)
```

Gets the value associated with the key `name` and returns it as a `double`.

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns The value, interpreted as a `double`. Null values evaluate to 0. Boolean values evaluate to 0 if false and 1 if true. Any other non-numeric value is converted to a `String` and then parsed as a `Double`.

```
public int getInt(String name)
```

Gets the value associated with the key `name` and returns it as an `int`.

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns	The value, interpreted as an <code>int</code> . Null values evaluate to 0. Boolean values evaluate to 0 if false and 1 if true. Any other non-numeric value is converted to a <code>String</code> and then parsed as an <code>Integer</code> .
----------------	--

```
public long getLong(String name)
```

Gets the value associated with the key `name` and returns it as a `long`.

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns	The value, interpreted as a <code>long</code> . Null values evaluate to 0. Boolean values evaluate to 0 if false and 1 if true. Any other non-numeric value is converted to a <code>String</code> and then parsed as a <code>Long</code> .
----------------	--

```
public Map getMap(String name)
```

Gets the value associated with the key `name` and returns it as an `IMu Map` object.

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns	The value, cast to a <code>Map</code> .
----------------	---

```
public Map[] getMaps(String name)
```

Gets the value associated with the key `name` and returns it as an array of `IMu Map` objects. This is a short-hand for:

```
getArray(name, Map[].class)
```

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns	The value, converted to a <code>Map[]</code> .
----------------	--

```
public String getString(String name)
```

Gets the value associated with the key `name` and returns it as a `String`.

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns	The value, interpreted as a <code>String</code> . Null values remain <code>null</code> . Any other non-string value is converted to a <code>String</code> using the object's <code>toString</code> method.
----------------	--

```
public String[] getStrings(String name)
```

Gets the value associated with the key `name` and returns it as an array of `Strings`. This is a short-hand for:

```
getArray(name, String[].class)
```

Parameters

<code>name</code>	The key whose associated value is to be returned.
-------------------	---

Returns

The value, converted to a `String[]`.

Class Module

`com.kesoftware.imu.Module`

Extends: `com.kesoftware.imu.Handler`

Provides access to an EMu module.

Constructors

```
public Module(String table, Session session)
```

Creates an object which can be used to access the EMu module specified by `table`.

Parameters

<code>table</code>	Name of the EMu module to be accessed.
<code>session</code>	A <code>Session</code> object to be used to communicate with the IMu server.

```
public Module(String table)
```

Same as `Module` above but a new session is created automatically using the `Session` class's default `host` and `port` values.

Properties

`String` `table`

getter `getTable()`

The name of the table associated with the `Module` object.

Methods

```
public int addFetchSet(String name, String columns)
```

Associates a set of columns with a logical name in the server. The name can be used instead of a column list when retrieving data using `fetch`.

Parameters

name	The logical name to associate with the set of columns.
columns	A string containing the names of the columns to be used when <code>name</code> is passed to <code>fetch</code> . The column names must be separated by a semi-colon or a comma.

Returns The number of sets (including this one) registered in the server.

Throws `IMuException` if a server-side error occurred.

```
public int addFetchSet(String name, String[] columns)
```

Same as `addFetchSet` above but the list of columns is passed as an array.

```
public int addFetchSet(String name, ArrayList<String> columns)
```

Same as `addFetchSet` above but the list of columns is passed as an array list.

```
public int addFetchSets(Map sets)
```

Associates several sets of columns with logical names in the server. This is the equivalent of calling `addFetchSet` for each entry in the map but is more efficient.

Parameters

sets	A <code>Map</code> containing a set of mappings between a name and a set of columns.
------	--

Returns The number of sets (including these ones) registered in the server.

Throws `IMuException` if a server-side error occurred.

```
public int addSearchAlias(String name, String columns)
```

Associates a set of columns with a logical name in the server. The name can be used when specifying search terms to be passed to `findTerms`. The search becomes the equivalent of an OR search involving the columns.

Parameters

name	The logical name to associate with the set of columns.
columns	A string containing the names of the columns to be used when <code>name</code> is passed to <code>findTerms</code> . The column names must be separated by a semi-colon or a comma.

Returns The number of aliases (including this one) registered in the server.

Throws `IMuException` if a server-side error occurred.

```
public int addSearchAlias(String name, String[] columns)
```

Same as `addSearchAlias` above but the list of columns is passed as an array.

```
public int addSearchAlias(String name, ArrayList<String> columns)
```

Same as `addSearchAlias` above but the list of columns is passed as an array list.

```
public int addSearchAliases(Map aliases)
```

Associates several sets of columns with logical names in the server. This is the equivalent of calling `addSearchAlias` for each entry in the map but is more efficient.

Parameters

<code>aliases</code>	A map containing a set of mappings between a name and a set of columns.
----------------------	---

Returns The number of sets (including these ones) registered in the server.

Throws `IMuException` if a server-side error occurred.

```
public int addSortSet(String name, String keys)
```

Associates a set of sort keys with a logical name in the server. The name can be used instead of a sort key list when sorting the current result set using `sort`.

Parameters

<code>name</code>	The logical name to associate with the set of columns.
-------------------	--

<code>keys</code>	A string containing the names of the keys to be used when <code>name</code> is passed to <code>sort</code> . The keys must be separated by a semi-colon or a comma.
-------------------	---

Returns The number of sets (including this one) registered in the server.

Throws `IMuException` if a server-side error occurred.

```
public int addSortSet(String name, String[] keys)
```

Same as `addSortSet` above but the list of keys is passed as an array.

```
public int addSortSet(String name, ArrayList<String> keys)
```

Same as `addSortSet` above but the list of keys is passed as an array list.


```
public int addSortSets(Map sets)
```

Associates several sets of sort keys with logical names in the server. This is the equivalent of calling `addSortSet` for each entry in the map but is more efficient.

Parameters

<code>sets</code>	A map containing a set of mappings between a name and a set of keys.
-------------------	--

Returns	The number of sets (including these ones) registered in the server.
----------------	---

Throws	<code>IMuException</code> if a server-side error occurred.
---------------	--

```
public ModuleFetchResult fetch(String flag, int offset, int count, String columns)
```

Fetches `count` records from the position described by a combination of `flag` and `offset`.

Parameters

<code>flag</code>	The position to start fetching records from. Must be one of:
-------------------	--

- "start"
- "current"
- "end"

<code>offset</code>	The position relative to <code>flag</code> to start fetching from.
---------------------	--

<code>count</code>	The number of records to fetch. A <code>count</code> of zero is permitted to change the location of the current record without returning any results. A <code>count</code> of less than zero causes all the remaining records in the result set to be returned.
--------------------	---

<code>columns</code>	A string containing the names of the columns to be returned for each record or the name of a column set which has been registered previously using <code>addFetchSet</code> . The column names must be separated by a semi-colon or a comma.
----------------------	--

Returns	A <code>ModuleFetchResult</code> object.
----------------	--

Throws	<code>IMuException</code> if a server-side error occurred.
---------------	--

```
public ModuleFetchResult fetch(String flag, int offset, int count, String[] columns)
```

Same as `fetch` above but the list of columns is passed as an array.

```
public ModuleFetchResult fetch(String flag, int offset, int count, ArrayList<String> columns)
```

Same as `fetch` above but the list of columns is passed as an array list.

```
public ModuleFetchResult fetch(String flag, int offset, int count)
```

Same as `fetch` above but no columns are requested. The results returned will still include the pseudo-column `rownum` for each fetched record.

```
public long findKey(long key)
```

Searches for a record with the key value `key`.

Parameters

`key` The key of the record being searched for.

Returns The number of records found. This will be either 1 if the record was found or 0 if not found.

Throws `IMuException` if a server-side error occurred.

```
public long findKeys(long[] keys)
```

Searches for records with key values in the array `keys`.

Parameters

`keys` The list of keys being searched for.

Returns The number of records found.

Throws `IMuException` if a server-side error occurred.

```
public long findKeys(ArrayList<Long> keys)
```

Same as `findKeys` above but the keys are passed in an array list.

```
public long findTerms(Terms terms)
```

Searches for records which match the search terms specified in `terms`.

Parameters

`terms` The search terms.

Returns An estimate of the number of records found.

Throws `IMuException` if a server-side error occurred.

```
public long findWhere(String where)
```

Searches for records which match the `TexQL` `where` clause.

Parameters

`where` The `TexQL` `where` clause to use.

Returns An estimate of the number of records found.

Throws `IMuException` if a server-side error occurred.

```
public long restoreFromFile(String file)
```

Restores a set of records from a file on the server machine which contains a list of keys, one per line.

Parameters

<code>file</code>	The file on the server machine containing the keys.
-------------------	---

Returns The number of records found.

Throws `IMuException` if a server-side error occurred.

```
public long restoreFromTemp(String file)
```

Restores a set of records from a temporary file on the server machine which contains a list of keys, one per line. Operates the same way as `restoreFromFile` except that the file parameter is relative to the server's temporary directory.

```
public long restoreFromTemp(String file)
```

Restores a set of records from a temporary file on the server machine which contains a list of keys, one per line. Operates the same way as `restoreFromFile` except that the file parameter is relative to the server's temporary directory.

Parameters

<code>file</code>	The file on the server machine containing the keys.
-------------------	---

Returns The number of records found.

Throws `IMuException` if a server-side error occurred.

```
public ModuleSortResult sort(String keys, String flags)
```

Sorts the current result set by the sort keys in `keys`. Each sort key is a column name optionally preceded by a "+" (for an ascending sort) or a "-" (for a descending sort).

Parameters

<code>keys</code>	A string containing the list of sort keys. The keys must be separated by a semi-colon or a comma.
<code>flags</code>	A string containing a set of flags specifying the behaviour of the sort. The flags must be separated by a semi-colon or a comma.

Returns A `ModuleSortResult` object. If the `report` flag has not been specified the result will be `null`.

Throws `IMuException` if a server-side error occurred.

```
public ModuleSortResult sort(String keys, String[] flags)
```

Same as `sort` above but the `flags` are passed as an array.

```
public ModuleSortResult sort(String keys, ArrayList<String> flags)
```

Same as `sort` above but the `flags` are passed as an array list.

```
public ModuleSortResult sort(String[] keys, String flags)
```

Same as sort above but the keys are passed as an array.

```
public ModuleSortResult sort(String[] keys, String[] flags)
```

Same as sort above but the keys and flags are passed as arrays.

```
public ModuleSortResult sort(String[] keys, ArrayList<String> flags)
```

Same as sort above but the keys are passed as an array and the flags are passed as an array list.

```
public ModuleSortResult sort(ArrayList<String> keys, String flags)
```

Same as sort above but the keys are passed as an array list.

```
public ModuleSortResult sort(ArrayList<String> keys, String[] flags)
```

Same as sort above but the keys are passed as an array list and the flags are passed as an array.

```
public ModuleSortResult sort(ArrayList<String> keys, ArrayList<String> flags)
```

Same as sort above but the keys and flags are passed as array lists.

Class ModuleFetchResult

`com.kesoftware.imu.ModuleFetchResult`

Provides results from a call to the `Module` `fetch` method.

Properties

`int` `count`

getter `getCount()`

The number of records returned in the result.

`long` `hits`

getter `getHits()`

The best estimate of the size of the result set after the `fetch` method has completed. When the `Module` object generates a result set using `findTerms` or `findWhere`, the number of matches is occasionally an overestimate of the true number of matches. After the `fetch` method has been called, the IMu server may have a better estimate of the true number of matches so it is included in the result.

`Map[]` `rows`

getter `getRows()`

The array of the records actually fetched. Each record is represented by a `Map` object, with the map keys being the names of the columns requested in the `fetch` call.

Class ModuleSortResult

`com.kesoftware.imu.ModuleSortResult`

Provides results from a call to the `Module sort` method. This is a recursive structure holding the information for one sort key. Information for secondary, tertiary and subsequent sort keys is stored in nested `ModuleSortResult` objects.

Properties

`int` `count`

getter `getCount()`

The number of distinct terms returned in the result.

`ModuleSortTerm[]` `terms`

getter `getTerms()`

The array of the distinct terms for a sort key. Each term is represented by a `ModuleSortTerm` object.

Class ModuleSortTerm

`com.kesoftware.imu.ModuleSortTerm`

Holds the information for a single distinct term in the results of a sort.

Properties

long count

getter `getCount()`

The number of occurrences of this term in the result set. For secondary or subsequent sort keys this is the number of occurrences for a given outer term.

ModuleSortResult nested

getter `getNested()`

Information regarding nested terms within this term. This will be `null` if there are no nested terms.

String value

getter `getValue()`

The value of the distinct term itself.

Class Session

`com.kesoftware.imu.Session`

Manages a connection to an IMu server. The server's host name and port can be specified in the constructor by setting properties on the object or by setting class-based default properties.

Class Properties

String `defaultHost`

getter `getDefaultHost()`

setter `setDefaultHost(String host)`

 The name of the host used to create a connection if no object-specific host has been supplied.

int `defaultPort`

getter `getDefaultPort()`

setter `setDefaultPort(int port)`

 The number of the port used to create a connection if no object-specific host has been supplied.

Constructors

`Session(String host, int port)`

Creates a `Session` object with the specified `host` and `port`.

`Session()`

Creates a `Session` object with the default host and port.

Properties

boolean	close	
	<i>getter</i>	getClose()
	<i>setter</i>	setClose(boolean close)
A flag controlling whether the connection to the server should be closed after the next request. This flag is passed to the server as part of the next request to allow it to clean up.		
String	context	
	<i>getter</i>	getContext()
	<i>setter</i>	setContext(String context)
The unique identifier assigned by the server to the current session.		
String	host	
	<i>getter</i>	getHost()
	<i>setter</i>	setHost(String host)
The name of the host used to create the connection. Setting this property after the connection has been established has no effect.		
int	port	
	<i>getter</i>	getPort()
	<i>setter</i>	setPort(int port)
The number of the port used to create the connection. Setting this property after the connection has been established has no effect.		
boolean	suspend	
	<i>getter</i>	getSuspend()
	<i>setter</i>	setSuspend(boolean suspend)
A flag controlling whether the server process handling this session should begin listening on a distinct, process-specific port to ensure a new session connects to the same server process. This is part of IMu's mechanism for maintaining state. If this flag is set to <code>true</code> , then after the next request is made to the server, the <code>Session's port</code> property will be altered to the process-specific port number.		

Methods

```
public void connect()
```

Opens a connection to an IMu server.

Throws `IMuException` if the connection could not be opened.

```
public void disconnect()
```

Closes the connection to the IMu server.

```
public void login(String user, String password, boolean spawn)
```

Logs in as the given user with the given password. If the `spawn` parameter is set to `true`, this will cause the server to create a new child process specifically to handle the newly logged in user's requests.

Parameters

<code>user</code>	The name of the user to login as.
<code>password</code>	The user's password for authentication.
<code>spawn</code>	A flag indicating whether the process should create a new child process specifically for handling the newly logged in user's requests.

Throws `IMuException` if the login request failed.

`Exception` (or another subclass) if a low-level socket communication error occurred.

```
public void login(String user, String password)
```

Same as `login` above except that the `spawn` parameter defaults to `true`.

```
public Map request(Map request)
```

Submits a low-level request to the IMu server.

Parameters

`request` A `Map` object containing the request parameters.

Returns A `Map` object containing the server's response.

Throw `IMuException` if a server-side error occurred.

Class TempInputStream

`com.kesoftware.imu.TempInputStream`

Extends: `java.io.FileInputStream`

This class is used to provide access to a temporary copy of a file returned from the server. This is most commonly used when a request is made to fetch a virtual multimedia column such as `resource`.

The class ensures that the temporary file that the stream is accessing is removed when the stream is closed or the stream object is finalised.

Constructors

`TempInputStream(File file)`

Creates an input stream to access the file referred to by `file`.

Note: This file is removed when the stream is closed or finalised.

Parameters

`file` Temporary file to be accessed.

Throws

`FileNotFoundException` if the file does not exist.

Methods

`public void close()`

Closes the input stream. The file associated with the stream is removed.

Throws

`IOException` if the stream access failed.

`public void finalize()`

Overrides the base class's `finalize` method. If `close` has not been called previously, the stream is closed and the file is removed.

Throws

`IOException` if the stream access failed.

Class Terms

`com.kesoftware.imu.Terms`

This class is used to create a set of search terms that is passed to the IMu server. A `Terms` object can be passed to the `findTerms` method of either a `Module` or `Modules` object.

Constructors

`Terms(TermsKind kind)`

Creates a new `Terms` object with the given `kind`. The `kind` can be either `TermsKind.AND` (for a set of AND terms) or `TermsKind.OR` (for a set of OR terms).

`Terms()`

Creates a new AND `Terms` object. This is the equivalent of:

`Terms(TermsKind.AND)`

Properties

`TermsKind` `kind`

getter

`getKind()`

The kind of terms list as specified when the object was constructed. Will be either:

- `TermsKind.AND`

-OR-

- `TermsKind.OR`

`Object[]` `list`

getter

`getList()`

The list of search terms themselves. Each element in the list can be either:

- A two or three element array comprising:
 - a column name
 - text to search for
 - an optional operator
- A nested `Terms` object

Methods

```
public void add(String name, String value, String operator)
```

Adds a new term to the list.

Parameters

name	The name of a column or a search alias.
value	The value to match.
operator	An operator to apply (such as "contains", "=", "<" etc.) for the server to apply when searching.

```
public void add(String name, String value)
```

Same as add above except no operator is specified. This is the preferred method for adding terms in many cases as it allows the server to choose the most suitable operator.

```
public Terms addAnd()
```

Adds an initially empty nested set of AND terms to the list. This is a shortcut for:

```
addTerms(TermsKind.AND)
```

Returns The newly added Terms object.

```
public Terms addOr()
```

Adds an initially empty nested set of OR terms to the list. This is a shortcut for:

```
addTerms(TermsKind.OR)
```

Returns The newly added Terms object.

```
public Terms addTerms(TermsKind kind)
```

Adds an initially empty nested set of terms to the list.

Returns The newly added Terms object.

Enum TermsKind

`com.kesoftware.imu.TermsKind`

An enumeration used to define the relationship between a set of terms in a `Terms` object.

Members

- | | |
|-----|---|
| AND | The relationship between the terms is AND |
| OR | The relationship between the terms is OR |

Index

A

A Simple Example • 30, 47
Accessing an EMu Module • 9
Attachments • 24

C

Class constants • 60
Class Handler • 57
Class IMu • 60
Class IMuException • 61
Class Map • 63
Class Module • 66
Class ModuleFetchResult • 73
Class ModuleSortResult • 74
Class ModuleSortTerm • 75
Class Properties • 76
Class Session • 76
Class TempInputStream • 79
Class Terms • 80
Column Sets • 29
columns • 20
Connecting to an IMu server • 7, 9
Constructors • 57, 61, 66, 76, 79, 80
count • 19

E

Enum TermsKind • 82
Example • 36, 38, 50
Examples • 14
Exceptions • 5, 55

F

findKey • 11
findKeys • 12
findTerms • 13
findWhere • 16
flag and offset • 18

flags • 34

G

Getting Information from Matching
Records • 17
Grouping a set of nested table columns • 27

H

Handlers • 7, 8

I

Introduction • 1

K

keys • 33

M

Maintaining State • 29, 47
Members • 82
Methods • 59, 62, 63, 67, 78, 79, 81
Multimedia • 41

N

Number of matches • 16

P

Properties • 58, 61, 66, 73, 74, 75, 77, 80

R

Reference • 17, 21, 57
Rename a Column • 26
Return Value • 37
Return Values • 21
Reverse Attachments • 25

S

Searching a Module • 10
Sorting • 33

T

Test Program • 4

U

Using IMu's Java library • 3