

Chapter 8

Practical Tile Storage

The previous chapter gave overviews for several different methods for storing tiled images. In this chapter we will present two fully-implemented techniques for storing tiled images together in large files. This type of method proved to be the best performing for writing, random reading, cached reading, and bulk copying. Furthermore, it is rather simple to implement. The first implementation shown is a fully functional method for writing and reading tile files and takes only about 200 lines of Java code for the reading, writing, and indexing methods.

Additionally, we will present the techniques with accompanying methods for creating tile indexes. These storage methods are designed to handle large and small sets of tiled images and are portable and updateable.

8.1 Introduction to Tile Indexes

Our goal is to store many hundreds or thousands of tiled images in a single file. This could be done by simply writing each image sequentially to a file. However, there would be no way to retrieve the images individually. There would be no way to know which tile address corresponded to which image. We could store the tile address before each image in our file, as shown in Figure 8.1. The problem with this method should be obvious. In order to access a specific tiled image, we have to scan the whole file. For tiled images sets of any significant size this method would be prohibitively inefficient. Instead we need to create a separate index into the file that will allow us to quickly look up the location of a specific tile in the file.

There are two principal ways we can construct the index:

- Sequential list of tile address to file position pairs
- Direct lookup table of file positions.

The simplest method is just to store the tile address, the position in the file and the size of the tiled image in a sequential list. This method is shown in Figure 8.2. The sequential list must still be searched for each tile query. However, the index

Record 1	Tile Address (Row, Column, Level)	Tile Image (N bytes)
Record 2	Tile Address (Row, Column, Level)	Tile Image (N bytes)
Record 3	Tile Address (Row, Column, Level)	Tile Image (N bytes)
Record 4	Tile Address (Row, Column, Level)	Tile Image (N bytes)

Fig. 8.1 Tile file with embedded addresses.

Index File

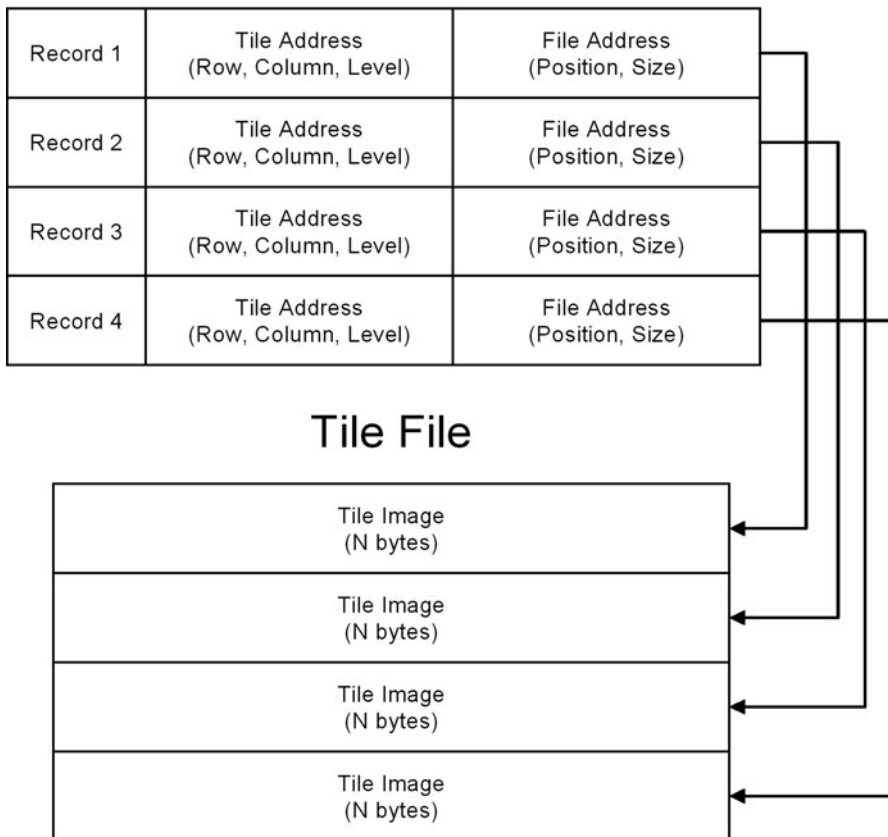


Fig. 8.2 Tile index stored as sequential list of addresses.

information is several orders of magnitude smaller than the actual tiles, thus scanning a separate index file for each query, while still inefficient, is must faster than scanning the whole file. An optimized variant of this method is to sort the data in the tile index. In this fashion a linear search could be used to speed up searching the index list.

The second method is to create a direct lookup table of file positions. For example, as shown in Figure 8.3, tile zoom level 3 has 8 columns and 4 rows, so the lookup table only requires 32 records. The table could be stored in a file in row-major or column-major format. The position of any given record can be directly computed, and only a single seek and read is required to retrieve the file position for a specific tile. We can store a null value in the lookup table to indicate that a tile does not exist for that specific address. If the table is stored in row-major order,

Lookup Table: Zoom Level 3

	Column 0	Column 1	Column 2	Column 3	Column 4	Column 5	Column 6	Column 7
Row 0	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)
Row 1	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)
Row 2	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)
Row 3	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)	File Address (Position, Size)

Fig. 8.3 Lookup table for zoom level 3.

Equation 8.1 is used to compute the position in the array of addresses.

$$p = j * C + i \tag{8.1}$$

where

i = column index

j = row index

C = number of columns

p = position of tile record

The disadvantage of this approach is that the size of our lookup table file grows by 4 times for each successive level. If the file address is stored as an 8 byte integer

and the size is stored as a 4 byte integer, we need 12 bytes for each record. Zoom level 17 contains 131,072 columns and 65,536 rows for a total of 8,589,934,592 tiles. This would require over 100 gigabytes just for the index file. If we had a tile set with a complete (or nearly complete) coverage of the earth's surface at that resolution, this approach would be appropriate.

However, this is unlikely. Most of the earth's surface is covered with water (liquid and ice) that is rarely imaged at high resolution. Few tile sets will cover even a fraction of the earth's surface. In these cases, we should develop an indexing method that provides direct lookup of tile locations, but also allows us to have lookup tables that cover only a subset of the entire level. This can be easily accomplished by providing for offsets attached to the index table. Rather than having all index tables start at (0,0) and covering the full range of tile addresses, we can provide external start and end addresses for index tables.

The next two sections will each present an algorithm for storing large amounts of tiled images. Each algorithm comes with its own unique method for indexing tiles. Those methods are modified versions of the direct lookup algorithm.

8.2 Storage by Zoom Level

Our first technique for storing tiles is to store all the tiles for a specific zoom level in a single file. This is the same technique that was tested and benchmarked in the previous chapter. This technique uses three files for each zoom level, one file for the tiled images and two files for the index.

The file containing the tiled images is simply a sequential list of tiled images. It first stores a magic number to serve as a sentinel value. Then it stores the tile's address and size. Finally it stores the tiled image data. The sentinel values and tile addresses are stored to make the tile images recoverable in the case that the tile file or index files become corrupted. Figure 8.4 shows the record structure for the tiled image file. Since the tiles do not have to be stored in any particular order, tiles can be written over a period of time. New tiles can be added to the file by simply writing them at the end of the file.

The index storage is slightly more complicated. Recall from the previous section that our lookup table based method can require a very large lookup table for the higher resolution zoom levels. To reduce the required size we have designed a two-step lookup table. We use the same approach to writing the lookup table from Figure 8.3, except that we only store rows in the index file that actually have tiles in them. So if our tile set only has 100 rows, then our tile index will only have 100 rows worth of tile addresses.

To accomplish this we have to create an additional index file, a row index file. This file contains a single value for each row in our set. If the row has any tiles, we store the location of that row's index records from the tile index file. If the row does not have any tiles, we store a null value in the file.

Magic Number	Tile Address (Row, Column)	Tile Size	Tile Image (N bytes)
Magic Number	Tile Address (Row, Column)	Tile Size	Tile Image (N bytes)
Magic Number	Tile Address (Row, Column)	Tile Size	Tile Image (N bytes)
Magic Number	Tile Address (Row, Column)	Tile Size	Tile Image (N bytes)

Fig. 8.4 Structure of tiled image file.

An example of this method is shown in Figure 8.5. We have used the same table from Figure 8.3, but we have assumed that rows 0 and 2 contain zero tiles. In this case, neither of those rows is stored in the index table, and the subsequent table is only half the size. Thus, the advantage of this technique is reduced space requirements for the index file. The disadvantage is that we have to do two seeks and reads to get the tile address. However, as shown in the previous chapter’s benchmarks, the performance is still very good.

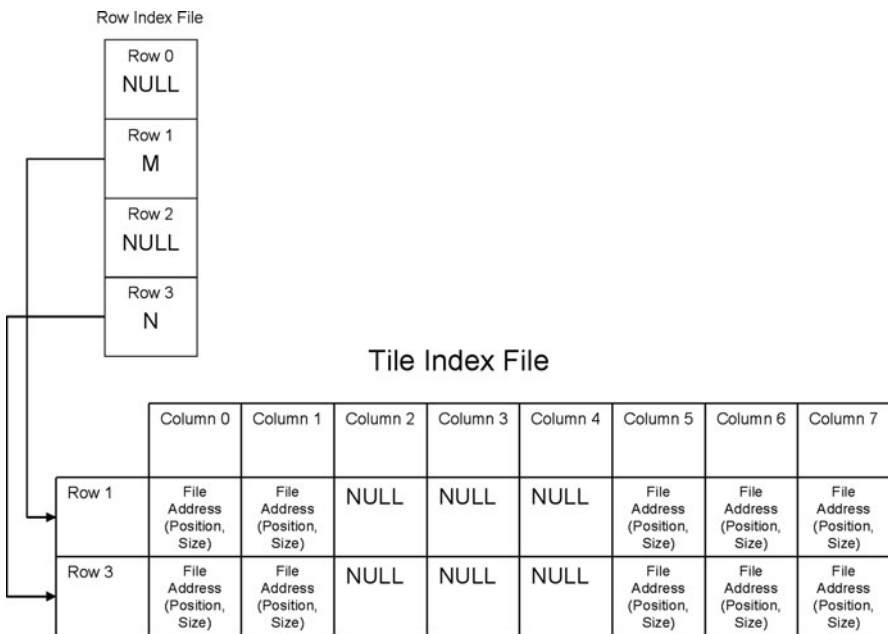


Fig. 8.5 Two-step tile index method.

To get the address for a specific tile, seek to the position of the row pointer in the row index file and read the value. If the value is non-null, use that value to position the tile index file. Then seek additional positions for the column index and read the tile address. Listings 8.1 and 8.2 present example code for writing and reading indexed tiles.

8.3 Introduction to Tile Clusters

The previous method works well and could be modified such that all levels can be contained in a single file. This would require addition of a third index file, a level index file similar to the row index file. Each tile address lookup would require 3 seeks and reads.

However, this method would not address two of the problems discussed in the tile creation chapter. Recall both the performance improvements made possible by caching tiles in memory (Section 6.1) and the requirement to have logically defined sub-groupings of tiles for distributed tile creation (Section 6.3.2). To address both of these requirements we propose a method for grouping tiled images in clusters.

Tiled image layers follow a pyramid type structure, see Figure 8.6. Each level has 4 times the number of tiles as its predecessor. Also, each lower resolution level is based on the image data from the next higher resolution level.



Fig. 8.6 Pyramid structure of tile images.

Our cluster-based grouping method starts by dividing the world into two clusters, (0,0) and (0,1). Figure 8.7 shows that division, and Figure 8.8 shows the structure of a cluster with 5 levels. The tiles that fall into the area marked by address (0,0) are stored in cluster (0,0), and all the tiles that fall into the area marked by address (0,1) are stored in cluster (0,1). By choosing this division we ensure that there are no tiles that overlap both clusters.

The number of tiles for a tile set with l levels is computed with Equation 8.2:

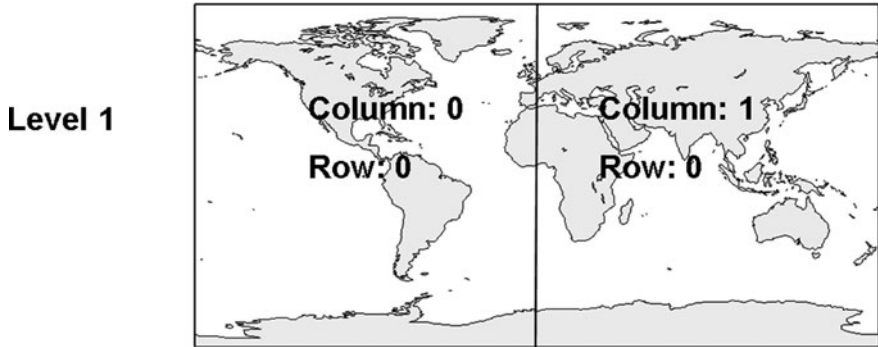


Fig. 8.7 World divided into two clusters.

Tile Cluster with 5 Levels




Fig. 8.8 Structure of a cluster with 5 levels.

$$N = \sum_{i=1}^L 2^i 2^{i-1} \tag{8.2}$$

The number of tiles for a cluster with l levels is the value from Equation 8.2 divided by two, or as shown in Equation 8.3:

$$N' = \sum_{i=1}^L 2^{2i-2} \tag{8.3}$$

8.4 Tile Cluster Files

To store tiles in cluster files, we must first set the number of levels to be stored. For a tile set with a base level of 7, we will need two cluster files, each with 7 levels of tiles and 5,461 tiles. Because the possible number of tiles is fixed for each cluster,

we can build a single fixed length lookup index and store it at the beginning of the cluster file. The index size will be the number of possible tiles times the size of the tile address record. After the index, we can store the tiled images sequentially in the file. Since we have an index, we do not need to store the tiles in any particular order. Figure 8.9 shows the file structure for a cluster file.




Fig. 8.9 Structure of a tile cluster file.

8.5 Multiple Levels of Clusters

When applying this method to tile sets with several more than 7 levels, we will experience the same problem discussed in Section 8.1. Our index will be too large. Imagine a tile set with only 100 tiles at level 15. Scaled versions of those 100 tiles will give us about 50 additional tiles with levels 14 to 1. That is a total of 150 tiles. If each tile is 50,000 bytes then the size of the tiles in total is 7.5 megabytes. However, a cluster file with 15 levels can have up to 357,913,941 tiles. If each index record takes up 12 bytes, the size of the index table would be 4,294 megabytes, or almost 600 times the size of the actual image data. This is a highly impractical consequence.

To alleviate this problem we allow multiple levels of clusters, with each level covering a continuous sub-range of levels. For example, if we have a tile set with 15 levels, we will have two levels of clusters, one level with contain tile levels 1-7, and the other level of clusters will contain levels 8-15. The first level contains 7 levels,

and the second level contains 8 levels. The indexes for multi-level cluster groups will never grow unmanageably large.

Continuing the example of a multi-level set of clusters, the first set, those with levels 1-7, can only have up to two clusters. While the second set, representing levels 8-15 can contain as many clusters as there are tiles in level 8. This number is 32,768. However, in practice we will only create clusters files when there are tiles that belong in the cluster. Few tile sets will have complete coverage of the whole world at a high resolution, and thus the full 32,768 would never actually be needed. The actual required number would fluctuate based on the size of the tile set.

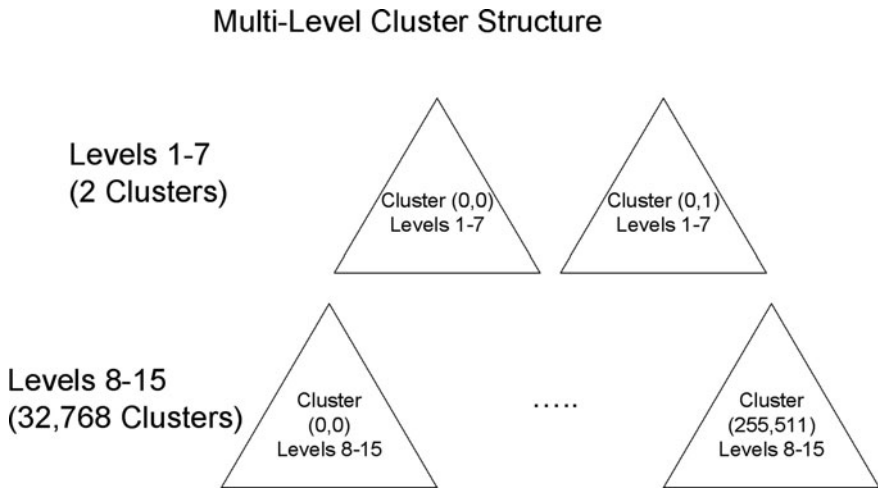


Fig. 8.10 Organization of multiple levels of tile clusters.

8.6 Practical Implementation of Tile Clusters

Listing 8.3 implements a cluster-based tile storage method. Since the internal file structure of our cluster files is relatively simple the implementation is relatively straightforward. The most difficult component of the practical implementation of our cluster-based storage system is the algorithm to determine in which cluster a given tile should be placed. That algorithm can be seen in the methods "getIndexPosition" and "getClusterFileTileAddress".

8.7 Application to Memory Cached Tiles

In Chapter 5, we saw potential performance improvements from holding tiles in memory while they were being created. The cluster-based storage technique works very well with this concept. To implement this with the cluster technique, we first ensure that our clusters are divided small enough to be held uncompressed completely in memory. If that is the case, then we alter our tile creation method to create tiles for one cluster at a time. We modify our clustered tile storage algorithm by simply adding a cache that holds all the tiles in memory as they are written. It writes them to file at the end of the tile creation process. This offers an additional performance improvement. If we write all the tiles at one time, we can write them sequentially and avoid using random file seeks, reads and writes. Random file accesses are generally slower than sequential accesses.

8.8 Application to Distributed Computing

The application of the tile clusters technique to distributed computing should be obvious. Multiple computer systems can be tasked with creating the tiles locally for specific cluster files. The individual cluster files can exist separately and function with minimal interaction, so they are a natural fit for highly distributed computing. After a cluster is completed, the single file can be copied back to a centralized location.

8.9 Performance Optimizations of Tile Cluster Method

There are several other potential performance optimizations available with the clustered storage technique. First, in our example code we opened and closed the various cluster files for each read and write. This is generally slower than maintaining constantly open files and reading and writing from them. Therefore, we might alter our algorithm to keep all the cluster files open throughout the process. However, many systems enforce a limit on the number of open files at any one time. Therefore, to get some performance benefit we can maintain a cache of recently opened files. The cache needs to be of sufficient size to ensure that open files are reused, but it must not be larger than the allowed maximum number of open files.

Since new tiles are written at the end of the file, this technique supports adding tiles over a period of time. When an existing tile is overwritten, the index is updated to point to the new tile. Old tiles remain in the file and take up space. Some developers may want to implement a system to try to re-use that space, either by trying to fit other tile images in the old space or by periodically rebuilding the entire cluster.

Finally, when tiles are served to users from clusters the performance can be quite good. Users typically view tiles for a specific area, and with our system those tiles

would be clustered in the same file. However, there is one case where the performance can be rather poor. Recall our example in which we had tiles in levels 1-15, and separated the clusters into groups of 1-7 and 8-15 levels. If a user is viewing level 8, and requesting several tiles from level 8, the system will have to access a different file for each tile. The benchmarks in the previous chapter showed that using a separate file per tile can be somewhat slow. A workaround to this problem is to build in some overlap in our cluster structure. Instead of a 1-7 and 8-15 break, we will use a 1-8 and 8-15 break. The tiles from level 8 are stored in two places. This does introduce some inefficiency; level 8 can have up to 32,768 tiles. But the read performance improvements may be worth the cost.

Listing 8.1 Output of indexed tiles by zoom level.

```

1  public class IndexedTileOutputStream {
2
3
4      static final long magicNumber = 0x772211ee;
5
6      String imagefilename;
7      String rowindexname;
8      String tileindexname;
9
10     RandomAccessFile imagefile;
11     RandomAccessFile rowindex;
12     RandomAccessFile tileindex;
13
14     long numRows, numcolumns;
15     int rowRecordSize = 8;
16     int tileRecordSize = 8 + 4;
17
18     public IndexedTileOutputStream(String folder, String setname, int level
19         ) {
20         imagefilename = folder + "/" + setname + "-" + level + ".tiles";
21         rowindexname = folder + "/" + setname + "-" + level + ".rowindex";
22         tileindexname = folder + "/" + setname + "-" + level + ".tileindex"
23         ;
24         numRows = TileStandards.zoomRows[level];
25         numcolumns = TileStandards.zoomColumns[level];
26
27         try {
28             imagefile = new RandomAccessFile(imagefilename, "rw");
29
30             //if the row file is empty, fill it with -1 to indicate empty
31             //values
32             rowindex = new RandomAccessFile(rowindexname, "rw");
33             if (rowindex.length() == 0) {
34                 rowindex.seek(0);
35                 for (int i = 0; i < numRows; i++) {
36                     rowindex.writeLong(-1L);
37                 }
38             }
39             tileindex = new RandomAccessFile(tileindexname, "rw");
40         } catch (Exception e) {
41             e.printStackTrace();
42         }
43     }
44
45     public void writeTile(long col, long row, byte[] data) {
46
47         try {
48             //position tile file to write at end of file
49             long writepos = imagefile.length();
50             imagefile.seek(writepos);
51
52             //write tile address and imagedata to file
53             //write two magic numbers so that tile records can be recovered
54             //in case of corrupted file
55             imagefile.writeLong(magicNumber);
56             imagefile.writeLong(magicNumber);
57             imagefile.writeLong(col);
58             imagefile.writeLong(row);
59             imagefile.writeInt(data.length);
60             imagefile.write(data);
61
62             //update index
63             updateIndex(col, row, writepos, data.length);
64
65         } catch (IOException e) {
66             e.printStackTrace();
67         }
68     }
69 }

```

```

63     }
64
65 }
66
67 private void updateIndex(long col, long row, long writepos, int length)
68 {
69     try {
70         //check if row is in the row index
71         long rowposition = rowRecordSize * row;
72         rowindex.seek(rowposition);
73         long rowpointer = rowindex.readLong();
74         if (rowpointer == -1L) {
75             //this means the row data is new and not already in the
76             //index
77             rowpointer = tileindex.length();
78             tileindex.seek(rowpointer);
79             //write an array of empty values
80             for (int i = 0; i < numcolumns; i++) {
81                 tileindex.writeLong(-1L);
82                 tileindex.writeInt(-1);
83             }
84             //write the position back to the original row index
85             rowindex.seek(rowposition);
86             rowindex.writeLong(rowpointer);
87         }
88         //compute offset into row for specific col
89         long offset = rowpointer + col * tileRecordSize;
90         //position tile index for writing the file address of the tile
91         //image
92         tileindex.seek(offset);
93         tileindex.writeLong(writepos);
94         tileindex.writeInt(length);
95     } catch (IOException e) {
96         e.printStackTrace();
97     }
98 }
99
100 public void close() {
101     try {
102         imagefile.close();
103         rowindex.close();
104         tileindex.close();
105     } catch (Exception e) {
106     }
107 }

```

Listing 8.2 Reading indexed tiles.

```

1 public class IndexedTileInputStream {
2
3     String imagefilename;
4     String rowindexname;
5     String tileindexname;
6
7     RandomAccessFile imagefile;
8     RandomAccessFile rowindex;
9     RandomAccessFile tileindex;
10
11     long numrows, numcolumns;
12     int rowRecordSize = 8;
13     int tileRecordSize = 8 + 4;
14
15 }

```

```

16 public IndexedTileInputStream(String folder, String setname, int level)
17 {
18     imagefilename = folder + "/" + setname + "-" + level + ".tiles";
19     rowindexname = folder + "/" + setname + "-" + level + ".rowindex";
20     tileindexname = folder + "/" + setname + "-" + level + ".tileindex"
21     ;
22     numRows = TileStandards.zoomRows[level];
23     numColumns = TileStandards.zoomColumns[level];
24
25     try {
26         imagefile = new RandomAccessFile(imagefilename, "rw");
27         rowindex = new RandomAccessFile(rowindexname, "rw");
28         tileindex = new RandomAccessFile(tileindexname, "rw");
29     } catch (Exception e) {
30         e.printStackTrace();
31     }
32
33     public byte[] getTile(long col, long row) {
34         try {
35             //check if row is in the row index
36             long rowposition = rowRecordSize * row;
37             rowindex.seek(rowposition);
38             long rowpointer = rowindex.readLong();
39             if (rowpointer == -1L) {
40                 //this means the row data is not in the index, and so the
41                 //tile doesn't exist
42                 return null;
43             }
44             //compute offset into row for specific col
45             long offset = rowpointer + col * tileRecordSize;
46             //position tile index for reading the position and size of the
47             //tile image
48             tileindex.seek(offset);
49             long tileposition = tileindex.readLong();
50             int size = tileindex.readInt();
51             if (tileposition == -1L) {
52                 //this means that the tile isn't there
53                 return null;
54             }
55             //adjust the tile position to skip the magic numbers and
56             //address information
57             long adjustedTilePosition = tileposition + 8 + 8 + 8 + 8 + 4;
58             byte[] data = new byte[size];
59             //position the image file and read the image data
60             imagefile.seek(adjustedTilePosition);
61             imagefile.readFully(data);
62             return data;
63         } catch (IOException e) {
64             e.printStackTrace();
65         }
66     }
67 }

```

Listing 8.3 Tile clusters implementation.

```

1 public class ClusteredTileStream {
2
3     static final long magicNumber = 0x772211ee;
4     private String location;
5     private String setname;
6     private int numlevels;
7     private int breakpoint;
8

```

```

9 |
10 |     public ClusteredTileStream(String location , String setname, int
11 |         numlevels , int breakpoint) {
12 |         this.location = location;
13 |         this.setname = setname;
14 |         this.numlevels = numlevels;
15 |         this.breakpoint = breakpoint;
16 |     }
17 |
18 |     public void writeTile(long row, long column, int level, byte[]
19 |         imagedata) {
20 |         //first determine the cluster that will hold the data
21 |         ClusterAddress ca = getClusterForTileAddress(row, column, level);
22 |         String clusterFile = getClusterFileForAddress(ca);
23 |         if (clusterFile == null) {
24 |             return;
25 |         }
26 |         File f = new File(clusterFile);
27 |
28 |         //if the file doesn't exist, set up an empty cluster file
29 |         if (!f.exists()) {
30 |             createNewClusterFile(f, ca.endlevel - ca.startlevel + 1);
31 |         }
32 |         try {
33 |             RandomAccessFile raf = new RandomAccessFile(f, "rw");
34 |
35 |             //write the tile and info at the end of the tile file
36 |             long tilePosition = raf.length();
37 |             raf.seek(tilePosition);
38 |             raf.writeLong(magicNumber);
39 |             raf.writeLong(magicNumber);
40 |             raf.writeLong(column);
41 |             raf.writeLong(row);
42 |             raf.writeInt(imagedata.length);
43 |             raf.write(imagedata);
44 |
45 |             //determine the position in the index of the tile address
46 |             long indexPosition = getIndexPosition(row, column, level);
47 |             raf.seek(indexPosition);
48 |
49 |             //write the tile position and size in the index
50 |             raf.writeLong(tilePosition);
51 |             raf.writeInt(imagedata.length);
52 |             raf.close();
53 |         } catch (Exception e) {
54 |             e.printStackTrace();
55 |         }
56 |     }
57 |
58 |     public byte[] readTile(long row, long column, int level) {
59 |         //first determine the cluster that will hold the data
60 |         ClusterAddress ca = getClusterForTileAddress(row, column, level);
61 |         String clusterFile = getClusterFileForAddress(ca);
62 |         if (clusterFile == null) {
63 |             return null;
64 |         }
65 |         File f = new File(clusterFile);
66 |
67 |         try {
68 |             RandomAccessFile raf = new RandomAccessFile(f, "r");
69 |
70 |             //determine the position in the index of the tile address
71 |             long indexPosition = getIndexPosition(row, column, level);
72 |             raf.seek(indexPosition);
73 |             long tilePosition = raf.readLong();
74 |             int tileSize = raf.readInt();

```

```

74         if (tilePosition == -1L) {
75             //tile is not in the cluster
76             raf.close();
77             return null;
78         }
79         byte[] imageData = new byte[tileSize];
80         //offset tile position for header information
81         long tilePositionOffset = tilePosition + 8 + 8 + 8 + 8 + 4;
82         raf.seek(tilePositionOffset);
83         raf.readFully(imageData);
84         raf.close();
85         return imageData;
86     } catch (Exception e) {
87         e.printStackTrace();
88     }
89     }
90     return null;
91 }
92 private long getIndexPosition(long row, long column, int level) {
93     ClusterAddress ca = this.getClusterForTileAddress(row, column,
94         level);
95     //compute the local address, that's the relative address of the
96     //tile in the cluster
97     int locallevel = level - ca.startlevel;
98     long localRow = (long) (row - (Math.pow(2, locallevel) * ca.row));
99     long localColumn = (long) (column - (Math.pow(2, locallevel) * ca.
100         column));
101     int numColumnsAtLocallevel = (int) Math.pow(2, locallevel);
102     long indexPosition = this.getCumulativeNumTiles(locallevel - 1) +
103         localRow * numColumnsAtLocallevel + localColumn;
104     //multiply index position times byte size of a tile address
105     indexPosition = indexPosition * (8 + 4);
106     return indexPosition;
107 }
108
109 public ClusterAddress getClusterForTileAddress(long row, long column,
110     int level) {
111     if (level > this.numlevels) {
112         //error, level is outside of ok range
113         return null;
114     }
115     int targetLevel = 0;
116     int endLevel = 0;
117     if (level < breakpoint) {
118         //tile goes in one of top two clusters
119         targetLevel = 1;
120         endLevel = breakpoint - 1;
121     } else {
122         //tile goes in bottom cluster
123         targetLevel = this.breakpoint;
124         endLevel = this.numlevels;
125     }
126     //compute the difference between the target cluster level and the
127     //tile level
128     int powerDiff = level - targetLevel;
129     //level factor is the number of tiles at level "level" for a
130     //cluster that starts at "target level"
131     double levelFactor = Math.pow(2, powerDiff);
132     // divide the row and column by the level factor to get the row and
133     //column address of the cluster we are using
134     long clusterRow = (int) Math.floor(row / levelFactor);
135     long clusterColumn = (int) Math.floor(column / levelFactor);
136     ClusterAddress ca = new ClusterAddress(clusterRow, clusterColumn,
137         targetLevel, endLevel);
138     return ca;
139 }

```



```

132     String getClusterFileForAddress(ClusterAddress ca) {
133         String filename = this.location + "/" + this.setname + "-" + ca.
            startlevel + "-" + ca.row + "-" + ca.column + ".cluster";
134         return filename;
135     }
136
137     //this methods create an empty file and fills the index with null
            values
138     void createNewClusterFile(File f, int numlevels) {
139         RandomAccessFile raf;
140         try {
141             raf = new RandomAccessFile(f, "rw");
142             raf.seek(0);
143             long tiles = this.getCumulativeNumTiles(numlevels);
144             for (long i = 0; i < tiles; i++) {
145                 raf.writeLong(-1L); //NULL position of tile
146                 raf.writeLong(-1L); //NULL size of tile
147             }
148             raf.close();
149         } catch (Exception e) {
150             e.printStackTrace();
151         }
152     }
153
154     public int getCumulativeNumTiles(int finallevel) {
155         int count = 0;
156         for (int i = 1; i <= finallevel; i++) {
157             count += (int) (Math.pow(2, 2 * i - 2));
158         }
159         return count;
160     }
161
162 }
163
164 public class ClusterAddress {
165
166     long row;
167     long column;
168     int startlevel;
169     int endlevel;
170
171     public ClusterAddress(long row, long column, int startlevel, int
        endlevel) {
172         this.row = row;
173         this.column = column;
174         this.startlevel = startlevel;
175         this.endlevel = endlevel;
176     }
177
178 }

```