

Tile-Based Geospatial Information Systems

John T. Sample • Elias Ioup

Tile-Based Geospatial Information Systems

Principles and Practices



Springer

John T. Sample
Naval Research Laboratory
1005 Balch Blvd.
Stennis Space Center, MS 39529
USA
john.sample@nrlssc.navy.mil

Elias Ioup
Naval Research Laboratory
1005 Balch Blvd.
Stennis Space Center, MS 39529
USA
elias.ioup@nrlssc.navy.mil

ISBN 978-1-4419-7630-7 e-ISBN 978-1-4419-7631-4
DOI 10.1007/978-1-4419-7631-4
Springer New York Dordrecht Heidelberg London

© Springer Science+Business Media, LLC 2010

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

To Nicole and Oliver
- John Sample

To Sarah and Georgette
- Elias Ioup

Preface

Tile-based mapping systems have grown to become the dominant form of mapping system with the rise of Web-based mapping tools. The origin of this book is a desire to collect all our discoveries, techniques, and best practices for creating a tiled-mapping system into one combined volume. The intent of this text is to provide a comprehensive guide to the theory behind creating a tiled-map system as well as a practical guide to create a concrete implementation.

Stennis Space Center, MS
May 2010

John Sample
Elias Ioup

Acknowledgements

The authors would like to thank the Naval Research Laboratory's Base Program, program element number 0602435N, for sponsoring this research. Additionally, the following people provided technical assistance without which this book would not have been possible: Perry Beason, Frank McCreedy, Norm Schoenhardt, Brett Hode, Bruce Lin, Annie Holladay, Juliette Ioup, and Hillary Mesick.

Contents

1	Introduction	1
1.1	Background of Web-Based Mapping Applications	1
1.2	Properties of tile-based mapping systems	2
1.3	Book Organization	2
2	Logical Tile Schemes	5
2.1	Introduction	5
2.2	Global Logical Tile Scheme	7
2.3	Blue Marble Example	10
2.4	Mercator-Based Schema	11
2.5	Variable Start Tile Schemes	12
2.6	Standardized Schema	15
	References	15
3	Tiled Mapping Clients	17
3.1	Tile Calculation	17
3.1.1	Discrete Map Scales	18
3.1.2	Continuous Map Scales	20
3.2	Tile Retrieval	22
3.2.1	Local Tile Storage	23
3.2.2	Network Tile Retrieval	23
3.3	Generating the Map View	25
3.3.1	Discrete Scales Map View	25
3.3.2	Continuous Scales Map View	26
3.4	Example Client	28
3.5	Survey of Tile Map Clients	28
4	Image Processing and Manipulation	35
4.1	Basic Image Concepts	35
4.2	Geospatial Images	37
4.2.1	Specialized File Formats	37

- 4.3 Image Manipulation 39
 - 4.3.1 Interpolation 1: Nearest Neighbor 44
 - 4.3.2 Interpolation 2: Bilinear 45
 - 4.3.3 Interpolation 3: Bicubic 46
- 4.4 Choosing Image Formats for Tiles 51
- 4.5 Choosing Tile Sizes 57
- 4.6 Tuning Image Compression 65
- References 79

- 5 Image Tile Creation 81**
 - 5.1 Tile Creation from Random Images 82
 - 5.2 Tile Creation Preliminaries 83
 - 5.2.1 Bottom-Up Tile Creation 83
 - 5.2.2 Choosing the Base Level for a Set of Source Images 83
 - 5.2.3 Pull-Based Versus Push-Based Tile Creation 87
 - 5.3 Tile Creation Algorithms 88
 - 5.3.1 Scaling Process for Lower Resolution Levels 89

- 6 Optimization of Tile Creation 97**
 - 6.1 Caching Tile Sets in Memory to Improve Performance 97
 - 6.2 Partial Reading of Source Images 99
 - 6.2.1 Reading Random Areas from Source Images 100
 - 6.2.2 Tile Creation with Partial Source Image Reading 103
 - 6.3 Tile Creation with Parallel Computing 103
 - 6.3.1 Multi-Threading of Tile Creation Algorithms 104
 - 6.3.2 Tile Creation for Distributed Computing 105
 - 6.4 Partial Updating of Existing Tiled Image Sets 108
 - References 116

- 7 Tile Storage 117**
 - 7.1 Introduction to Tile Storage 117
 - 7.2 Storing Image Tiles as Separate Files 118
 - 7.3 Database-Based Tile Storage 121
 - 7.4 Custom File Formats 121
 - 7.5 Comparative Performance 122
 - 7.5.1 Writing Tests 123
 - 7.5.2 Reading Tests 124
 - 7.6 Storage of Tile Metadata 126
 - 7.7 Storage of Tiles in Multi-Resolution Image Formats 126
 - 7.8 Memory-Cached Tile Storage 127
 - 7.9 Online Tile Storage 127

- 8 Practical Tile Storage** 133
 - 8.1 Introduction to Tile Indexes 133
 - 8.2 Storage by Zoom Level..... 136
 - 8.3 Introduction to Tile Clusters..... 138
 - 8.4 Tile Cluster Files 139
 - 8.5 Multiple Levels of Clusters 140
 - 8.6 Practical Implementation of Tile Clusters..... 141
 - 8.7 Application to Memory Cached Tiles 142
 - 8.8 Application to Distributed Computing 142
 - 8.9 Performance Optimizations of Tile Cluster Method 142

- 9 Tile Serving** 151
 - 9.1 Basics of HTTP 151
 - 9.2 Basic Tile Serving 152
 - 9.3 Tile Serving Scheme with Encoded Parameters..... 153
 - 9.4 Tile Serving Scheme with Encoded Paths 155
 - 9.5 Service Metadata Alternatives 156
 - 9.6 Conclusions 157
 - References 164

- 10 Map Projections** 165
 - 10.1 Introduction to Datums, Coordinate Systems, and Projections..... 165
 - 10.1.1 The Shape of the Earth..... 165
 - 10.1.2 Datums 166
 - 10.1.3 Coordinate Systems 169
 - 10.2 Map Projections 169
 - 10.2.1 Different Map Projections 170
 - 10.2.2 Cylindrical Equidistant Projection 171
 - 10.2.3 Cylindrical Equal-Area Projection 172
 - 10.2.4 Mercator 172
 - 10.2.5 Universal Transverse Mercator 172
 - 10.3 Point Reprojection..... 175
 - 10.4 Map Reprojection 177
 - 10.4.1 Affine Transforms 177
 - 10.4.2 Interpolation 179
 - 10.4.3 Point-wise Reprojection..... 180
 - 10.4.4 Tablular Point-Wise Reprojection..... 182
 - 10.5 Map Projections for Tiled Imagery 184
 - 10.5.1 Storing Tiles in the Geodetic Projection 184
 - 10.5.2 Storing Tiles in the Mercator Projection 185
 - 10.5.3 Other Projections 186
 - 10.5.4 Which Projection for a Tiled-Mapping System? 187
 - 10.6 Conclusion 188
 - References 191

- 11 Tile Creation using Vector Data** 193
 - 11.1 Vector Data 193
 - 11.2 Tile Creation 194
 - 11.3 Queries 196
 - 11.4 Storage 196
 - 11.4.1 Database Storage 197
 - 11.4.2 File System Storage 200

- 12 Case Study: Tiles from Blue Marble Imagery** 205
 - 12.1 Pull-Based Tiling 205
 - 12.2 Push-Based Tiling 207
 - 12.3 Results 207

- 13 Case Study: Supporting Multiple Tile Clients** 221
 - 13.1 KML Server 221
 - 13.1.1 Static KML Example 221
 - 13.1.2 Dynamic KML Example 223
 - 13.2 WMS Server 223
 - 13.2.1 WMS Servlet Implementation 224
 - References 233

- Index** 235

Chapter 1

Introduction

This book is intended to provide the reader with a thorough understanding of the purpose and function of tile-based mapping systems. In addition, it is meant to be a technical guide to the development of tile-based mapping systems. Complex issues like tile rendering, storage, and indexing are covered along with map projections, network communication, and client/server applications. Computer code as well as numerous mathematical formulae are included to provide the reader with usable forms of the algorithms presented in this book.

1.1 Background of Web-Based Mapping Applications

The first Web-based mapping applications were introduced in the mid to late 1990's. They included Yahoo! Maps, MapQuest, and Microsoft's TerraServer. These providers offered mapping applications through a Web browser. Their map navigation systems were rudimentary. Some allowed simple map movements by requiring users to click on navigation arrow buttons surrounding the map view. When users clicked on an arrow, the map moved a predetermined amount in the direction clicked. There were also buttons for zooming in and out. Others allowed users to drag and draw boxes on the map to relocate the map view.

All of these systems had several disadvantages, including slow rendering and downloading of map views because the map view was often represented by a single large image file. Each time the map was moved to the left or right; the entire image would be re-rendered and re-sent to the client even though only a portion of the image was new. However, the interfaces were relatively simple and had several advantages to developers. Basic interfaces were well suited to early Web browsers. The map interface could be written entirely in HTML or with very minimal JavaScript. Second, since all navigations were fixed, map servers could cache rendered maps. Other map viewers adopted a map view and navigation style more similar to desktop GIS systems. These systems were more complicated and used browser plugin technology and development platforms like Java or Flash.

Google Maps was introduced in 2005 and dramatically changed the way people viewed maps. Instead of clunky and slow map navigation methods, Google Maps provided what has come to be known as a "Slippy Map" type interface. That interface allowed users to quickly move and zoom the map and yet was written entirely in HTML and JavaScript. Soon many more Web mapping applications appeared with a similar style map interface. Eventually slippy map type interfaces appeared in many places including portable computing devices and cell phones.

A key enabling technology behind this new generation of mapping applications was the concept of tile-based mapping. Mapping applications were made responsive by using background maps that had been broken into smaller tiled images. Those tiles were stored, already rendered, on a central server. Because they were already rendered, they could be sent to clients quickly. The tiles were discretely addressed so they could be cached by Internet caching services and by clients' own browsers. The map images were broken into smaller pieces, so when users navigated the map view, only the new parts of the map had to be resent from the server.

1.2 Properties of tile-based mapping systems

Tile-based mapping systems have several core properties which distinguish them from other types of mapping systems. We have defined what we believe to be those core properties, and they are as follows:

1. Map views are based on multiple discrete zoom levels, each corresponding to a fixed map scale.
2. Multiple image tiles are used to virtualize a single map view.
3. Image tiles are accessible using a discrete addressing scheme.
4. Tiled images stored on a server system are sent to the client with minimal processing; as much processing is done ahead of time as is possible.

The following are important but optional properties of tile-based mapping systems.

1. Tile addressing follows a single global projection.
2. Tiles are primarily distributed using a client/server system architecture.
3. Tiles are organized into relatively few, fixed layers.

1.3 Book Organization

This book is organized to take the reader through the logical development of a complete tile-based mapping system with small detours into important topics along the way. Chapter 2 introduces logical tile addressing schemes that any tile system must implement. It discusses some common schemes used by popular Web mapping systems and defines the common tile scheme that will be used throughout this book.

Chapter 3 gives an overview of the challenges and the techniques used to overcome these challenges to develop client software for tile-based mapping systems. An example client application is shown with source code. Chapter 4 provides extensive background into techniques needed to process source data images into tiled images. Chapters 5 and 6 provide a detailed look at techniques for creating sets of tiled images. Chapters 7 and 8 explain how to efficiently store, index, and retrieve tiled images. Several techniques are detailed, implemented, and benchmarked. Chapter 9 shows the reader how to create a Web based server for tiled images. Chapter 10 introduces and explains map projections within the context of tile based mapping. Chapter 11 explains how vector mapping data can be used in a tile-based environment. Finally, Chapters 12 and 13 are detailed case studies of real-world usage of the techniques presented in this book.

Most chapters include computer code listings in Java and Python. Java and Python are two of the most commonly used programming languages for geospatial programming. Short code segments are interspersed with the chapter text, while longer code segments are placed at the end of each chapter. The code sections are intended to provide readers with example implementations of the algorithms explained in the book.

Chapter 2

Logical Tile Schemes

2.1 Introduction

Tile-based mapping systems use a logical tile scheme that maps positions on the Earth to a two-dimensional surface and divides that surface into a series of regularly spaced grids (see Figure 2.1). The logical tile scheme defines the discrete addressing of map tiles, the method for generating multiple zoom levels of tiles, and the translation method between tile addresses and a continuous geospatial coordinate system.

The logical tile scheme is the foundational element of a tile-based mapping system. It is a multi-resolution, regularly spaced grid. Each scheme is typically tied to a single two-dimensional map projection (for more on map projections see Chapter 9). This addressing scheme allows a tiled image to be accessed directly with discrete coordinates. For example, instead of requesting a map image with a bounding rectangle delineated with continuous real numbers like $[-100.0, 30.0]$ to $[-80.0, 40.0]$, a tile can be requested from a grid with level, column, and row addresses delineated with discrete integer values.

The logical tile scheme consists of a mapping between the address of a tile to the geospatial coordinates for the area covered by the tile. In general, there are several ways to develop a logical tile scheme. We could make custom schemes that match the bounds and dimensions of each individual data set, or we could create a single global tile scheme that can be applied to all data sets.

Each method has its benefits. In developing a logical tile scheme, we have to choose a series of image pixel resolutions, one for each level. If we can develop a new scheme for each data set, then we can choose image pixel resolutions that exactly match the resolution of our data set. Using a global common scheme, we have to use the predefined resolutions. Pre-defined resolutions will force us to rescale our source images to match. If we level down, we sacrifice some native resolution, and if we level up, we are using more storage space than is needed. However, if we use a custom scheme for each different data set, we will have interoperability issues in combining the data sets. Also, client software systems will have difficulty using

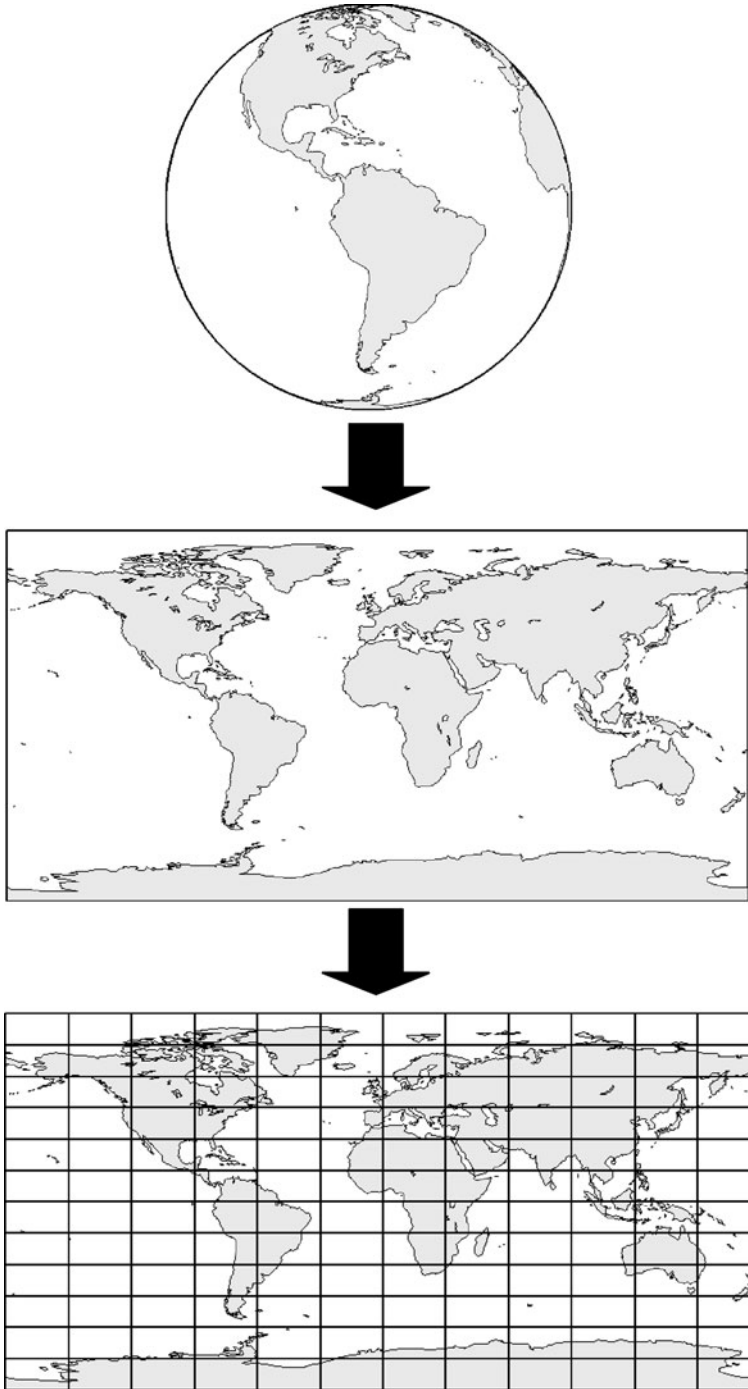


Fig. 2.1 Mapping from a spherical Earth to a two-dimensional surface to a gridded surface.

tiles from different schemes with different resolutions. For our purposes, we choose to use a common global tile scheme that is the same across all data sets. This choice sacrifices flexibility but simplifies system development and use.

2.2 Global Logical Tile Scheme

The global logical tile scheme presented in this book has been developed to be easy to understand and implement. We start with the geodetic projection, which simply portrays the Earth as a rectangle 360 degrees wide and 180 degrees tall. Our base projection has a natural 2-to-1 aspect ratio, and so does our logical tile scheme. At zoom level 1, our tile scheme has 1 row and 2 columns (Figure 2.2).




Fig. 2.2 Global tile scheme at zoom level 1.

For each subsequent level, we double the number of rows and columns. Since we have doubled each dimension, each subsequent level has 4 times the number of tiles as the previous level. As we increase zoom levels, each tile is divided into four sub-tiles (Figure 2.3).

We can continue this process and define as many levels as are needed. In practice, 20 levels are sufficient for almost any mapping data available. To simplify the mathematics, we start our indexes at 0 instead of 1. Our scheme can be completely defined mathematically. Equation 2.1 gives the number of columns for a given level i . Equation 2.2 gives the number of rows for a given level i .

$$C_i = 2^i \tag{2.1}$$

$$R_i = 2^{i-1} \tag{2.2}$$

Equations (2.3) through (2.6) relate a tile's address back to a geographic bounding rectangle.




Fig. 2.3 Global tile scheme at zoom level 2.

$$\lambda_{min} = c \frac{360.0}{2^i} - 180.0 \quad (2.3)$$

$$\lambda_{max} = (c + 1) \frac{360.0}{2^i} - 180.0 \quad (2.4)$$

$$\phi_{min} = r \frac{180.0}{2^{i-1}} - 90.0 \quad (2.5)$$

$$\phi_{max} = (r + 1) \frac{180.0}{2^{i-1}} - 90.0 \quad (2.6)$$

where

$c \equiv$ column

$r \equiv$ row

$\lambda \equiv$ longitude

$\phi \equiv$ latitude

$i \equiv$ zoom level (2.7)

Chapter 4 discusses choosing the pixel dimensions of tiled images. Once pixel dimensions are chosen, we can compute the resolution of our tiled images in terms of degrees per pixel (DPP). DPP is useful for relating tiled image zoom levels to continuous zoom levels used by many mapping applications. Equation (2.8) is used to calculate degrees per pixel.

$$DPP = \frac{360.0}{2^i} p \quad (2.8)$$

where

$p \equiv$ number of pixels per tile

$i \equiv$ zoom level

Zoom Level	Number of Columns	Number of Rows	Number of Tiles	Degrees Per Pixel
1	2	1	2	0.3515625000
2	4	2	8	0.1757812500
3	8	4	32	0.0878906250
4	16	8	128	0.0439453125
5	32	16	512	0.0219726563
6	64	32	2048	0.0109863281
7	128	64	8192	0.0054931641
8	256	128	32768	0.0027465820
9	512	256	131072	0.0013732910
10	1024	512	524288	0.0006866455
11	2048	1024	2097152	0.0003433228
12	4096	2048	8388608	0.0001716614
13	8192	4096	33554432	0.0000858307
14	16384	8192	134217728	0.0000429153
15	32768	16384	536870912	0.0000214577
16	65536	32768	2147483648	0.0000107288
17	131072	65536	8589934592	0.0000053644
18	262144	131072	34359738368	0.0000026822
19	524288	262144	137438953472	0.0000013411
20	1048576	524288	549755813888	0.0000006706

Table 2.1 The number of rows, columns, and tiles as well as the degrees per pixel for zoom levels 1 through 20 (assuming 512x512 pixel tiles).

Equations (2.9) and (2.10) show the method for locating the tile that contains a specific geographic coordinate, given a zoom level.

$$c = \lfloor (\lambda + 180.0) * \frac{360.0}{2^i} \rfloor \quad (2.9)$$

$$r = \lfloor (\phi + 90.0) * \frac{180.0}{2^{i-1}} \rfloor \quad (2.10)$$

where

$c \equiv$ horizontal tile index

$r \equiv$ vertical tile index

$\lambda \equiv$ longitude

$\phi \equiv$ latitude

$i \equiv$ zoom level of map view

2.3 Blue Marble Example

To better illustrate the concept, let us begin our first example. We want to define a logical time scheme suited to serving NASA's Blue Marble¹ imagery as tiles. It is satellite derived imagery of the whole earth and provides a good example data set for this book.




Fig. 2.4 A Blue Marble image.

For our example, we start with a single JPEG image that is 4096 pixels wide and 2048 pixels high. The image covers the entire earth and thus has a bounding rectangle of $(-180, -90)$ to $(180, 90)$.

For our example, we are going to use 512x512 pixel tile images. (Chapter 4: Image Processing and Manipulation for GIS will discuss how to choose the proper tile image size for a given application.) Dividing our image width (4096) by our tile width (512) gives us an even 8 tiles across. Likewise, we get an even 4 tiles vertically. From Table 2.1, this is exactly equivalent to zoom level 3. Therefore, we can use zoom level 3 as the base level for our example tile scheme.

In the previous example, our source image matched nicely with our global tile scheme. However, many data sets will not. Suppose we have a single image covering a small geographic area, $(-91.5, 30.2)$ to $(-91.4, 30.3)$, and the image is 1000 by 1000 pixels in size. The image covers a square 0.1 degrees by 0.1 degrees. The resolution of Level 1 is 0.35156. At that resolution, our entire image would only take up 0.28 by 0.28 of a pixel or 7.84% of a pixel. In other words, it would hardly be visible at Level 1. The DPP resolution of the image is $0.1/1000$, or 0.0001. This falls between levels 12 and 13 of our global tile scheme. In this case, it might be better to create a custom tile scheme. A method for defining custom schemes is presented in Section 2.5.

¹ <http://earthobservatory.nasa.gov/Features/BlueMarble>

2.4 Mercator-Based Schema

Throughout this book we will focus primarily on tiling systems and data that use the simple Plate Carrée projection, which is also known as the geographic projection. This projection is straightforward to work with and gives us a two-dimensional representation of the earth with a 2-to-1 horizontal to vertical aspect ratio. However, the geographic projection has several shortcomings. At high latitudes, shapes and angles become distorted. To avoid this distortion, many tiling systems use a different base projection for their tiling schemes. The spherical Mercator projection is used by Google Maps, Microsoft Bing Maps, and Yahoo! Maps. Chapter 9 will discuss the details of the Mercator projection. For the purposes of defining a tiling scheme, this projection is significant because it yields a global two-dimensional representation of the earth with a 1 to 1 aspect ratio (see Figure 2.5).



Fig. 2.5 Mercator projection.

Google, Microsoft, and Yahoo! all use a global image similar to what is shown in Figure 2.5 as the top level image in their tiling schemes. Higher resolution zoom levels are generated by dividing each tile into 4 sub-tiles. The only significant difference between these three schemes is in their respective methods for addressing and numbering the tiles. Google Maps uses a simple pair of coordinates to address tiles for a specific zoom level. They set the origin at the top, left of the map. Figure 2.6 shows the addressing for Google Maps at their zoom level 1.

Microsoft's Bing Maps also uses the top-left for its origin but uses a sequential numbering scheme, as shown in Figure 2.7. As the zoom level increases, each tile is divided into 4 sub-tiles. The sub-tiles are sequentially numbered 0 to 3, and that number is concatenated to the number of the parent tile to form the address of the




Fig. 2.6 Google Maps tile addressing at zoom level 1.

sub-tiles. Tile 0 is divided into sub-tiles 00, 01, 02, and 03 as shown in Figure 2.8. So, a tile at the 17th zoom level 17 would have 17 digits, one for each zoom level. This numbering scheme makes computing the addresses of sub-tiles trivial. However, relating tile addresses to geographic coordinates, and vice-versa, will require much more computation than the other methods of addressing tiles.

2.5 Variable Start Tile Schemes

NASA World Wind is a freely available virtual globe software system. It provides native support for tiled image sets as map backgrounds on the globe. The default World Wind tile system is very similar to the logical tile scheme we presented in Section 2.1. Like our scheme, World Wind uses the geographic projection. It also uses the bottom-left of the earth as its origin for tile addressing. It differs from our system in that rather than a 2 x 1 tile matrix for its first zoom level, it uses a 10 x 5 matrix as shown in Figure 2.10.

Yahoo! Maps uses a method very similar to Google Maps, except they set their origin tile at the left, middle of the earth, see Figure 2.9.

The World Wind system also allows tiled data sets that do not start at the global level. They use a concept called "Level Zero Tile Size" (LZTS) to define the dimensions of a tile at the lowest resolution level. Then they define a start scale to set the custom tile set's start point. Tiles at the zero level of their default tile scheme are 36.0 degrees by 36.0 degrees, so the LZTS for that scheme is 36.0. The start point