

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

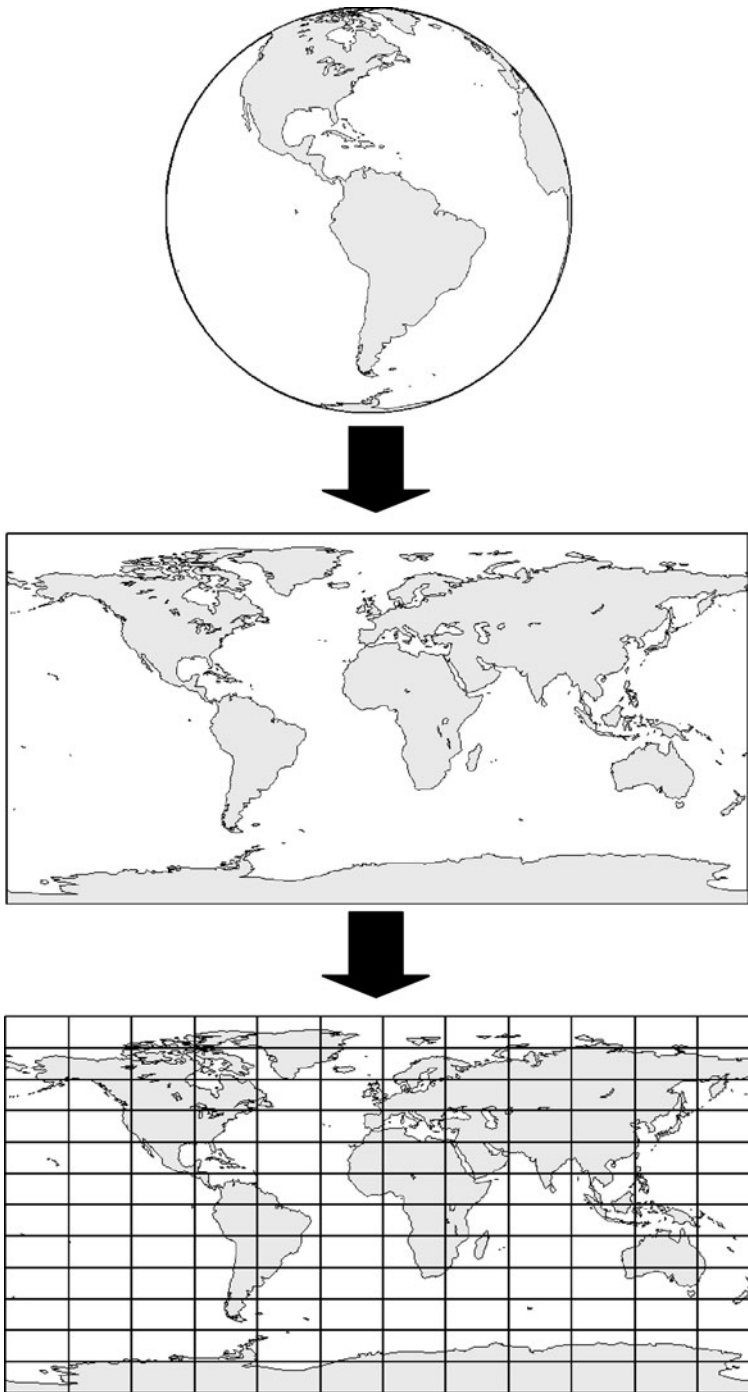


Fig. 2.1 Mapping from a spherical Earth to a two-dimensional surface to a gridded surface.

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.

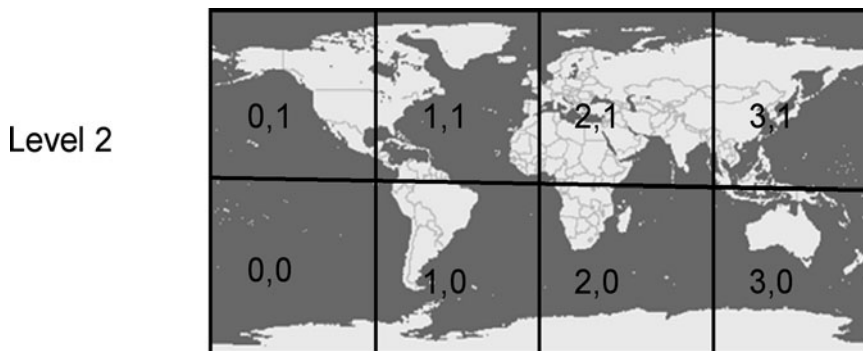


Fig. 2.3 Global tile scheme at zoom level 2.

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

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

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

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

where

$$\begin{aligned} c &\equiv \text{column} \\ r &\equiv \text{row} \\ \lambda &\equiv \text{longitude} \\ \phi &\equiv \text{latitude} \\ i &\equiv \text{zoom level} \end{aligned} \quad (2.7)$$

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.

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

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

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




Fig. 2.7 Microsoft Bing Maps tile addressing at zoom level 1.




Fig. 2.8 Subtile addressing in Microsoft Bing Maps.




Fig. 2.9 Yahoo! Maps tile addressing at zoom level 1.

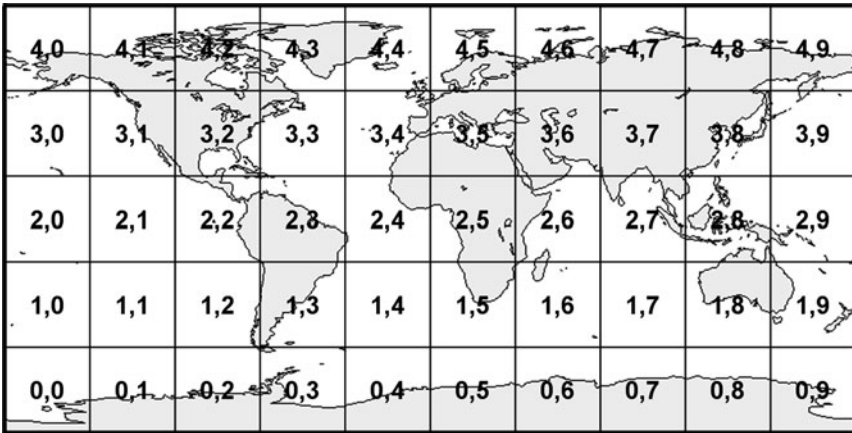


Fig. 2.10 NASA World Wind tile scheme at top zoom level.

for that scheme is at level 0. One could define a tile set with a different LZTS that did not start until a later zoom level, such as 1, 5, or 10. Because of the way World Wind renders the tiled images, custom LZTS values should be a factor of 180.

2.6 Standardized Schema

There are several efforts underway to standardize tile schemes and the way they are communicated. The Web Map Tile Service (WMTS) standard was recently finalized by the Open Geospatial Consortium [1]. It provides a standard but flexible way of defining the capabilities of a tile service and how to interface with it. WMTS does not require the use of one specific tile scheme, resolution set, or projection. Instead, it provides a standard means of defining these properties so clients and servers may be connected together. The WMTS standard does address tiles using matrix coordinates; the top-left tile is addressed as (0,0). However, other properties of the tile scheme are left to the service creator. Multiple different projections are allowed, including the Geodetic and Mercator projections. No restriction is made on which tile scales are made available, only that they be defined using a map scale, meaning the ratio of a distance on the map to a distance on the ground. The map scale is intended only as an identifier for a given zoom level, since it is accurate only near the equator. Tile size may vary over scale, and there may be no relation between the tile matrix dimensions and the scales. Of course, allowing this level of flexibility increases the difficulty of writing a generic client to support a generic WMTS server. To reduce this complexity, the WMTS standard supports a set of well known scale sets that a server may support. By implementing a well known scale set, the server becomes compatible with a wider range of clients. The set of scales in our tile scheme and the set of scales in the Google Maps Mercator tile scheme are both included WMTS well known scale sets. WMTS supports Key-Value-Pair, RESTful, and SOAP request formats for accessing tiles.

Another attempt to create a tile service standard is the Tile Map Service (TMS) specification [2]. The TMS specification is not backed by a standards body but has achieved some level of common usage with a number of servers and clients. It is similar to the WMTS standard in that it allows multiple different tile schemes to be specified. The TMS specification allows the use of arbitrary scales defined by units per pixel. The origin tile may be specified by the server unlike in WMTS where it is always the top-left tile. The tile size may be specified as well. As with WMTS, it supports profiles that specify a map scale and map projection. Both our Geodetic tile scheme and Google Maps Mercator tile scheme are supported profiles. The TMS format supports only a RESTful URL request for tiles.

References

1. Joan Masó, K.P., Julia, N.: OpenGIS Web Map Tile Service Implementation Standard. Open Geospatial Consortium Specification (2010)
2. Ramsey, P.: Tile map service specification. URL http://wiki.osgeo.org/wiki/Tile_Map_Service_Specification