Mastering ArcGIS

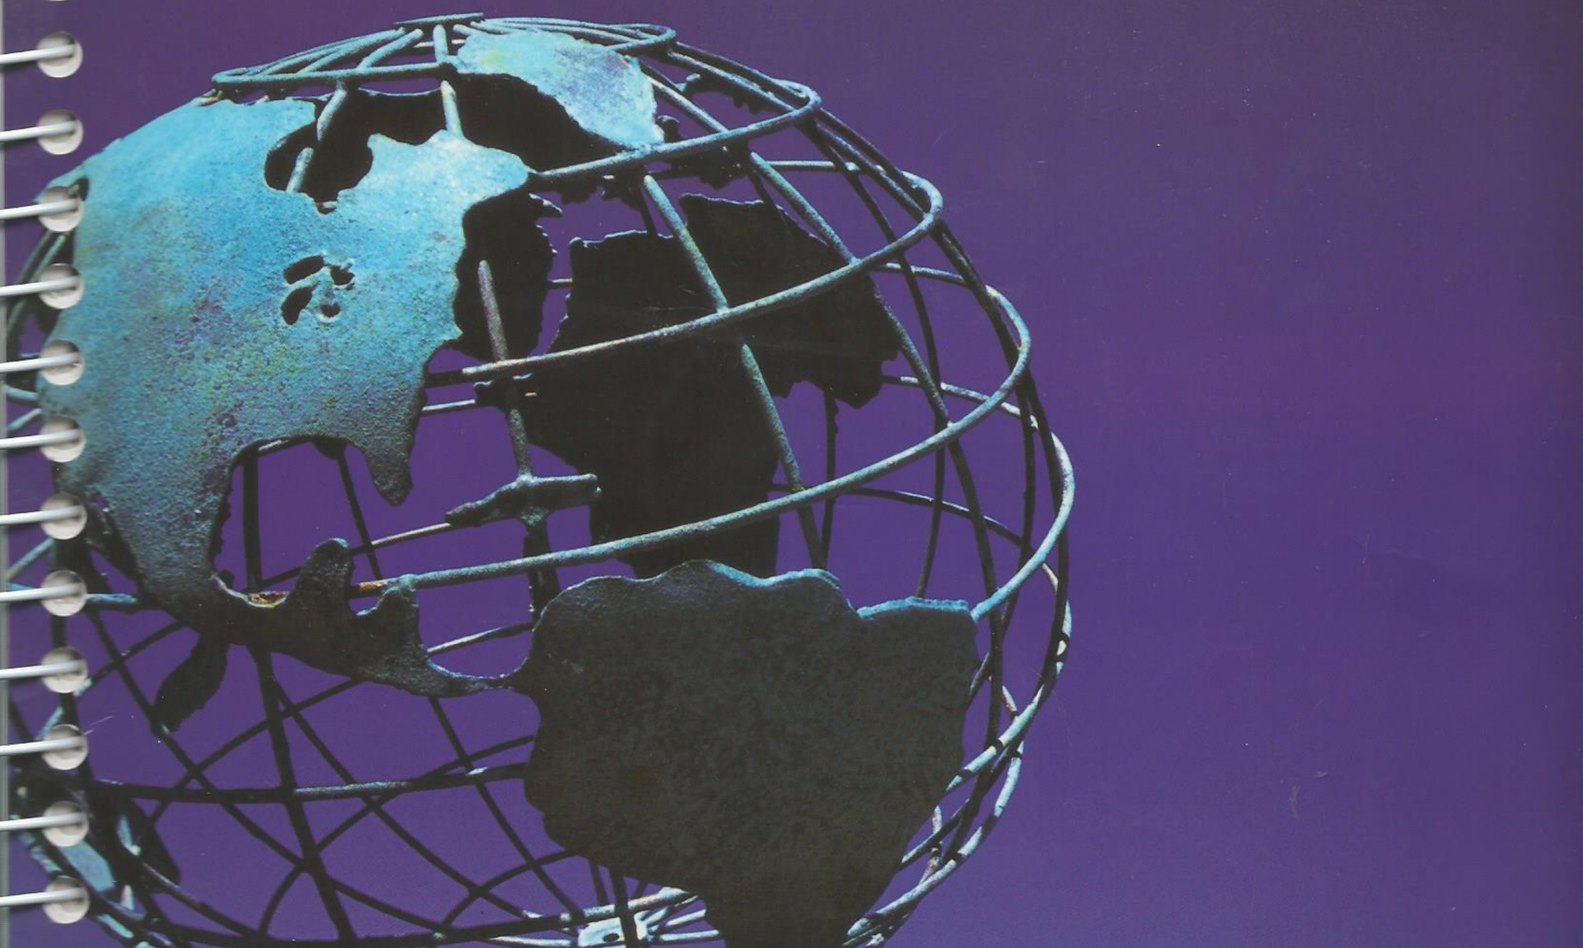

Sixth Edition

Maribeth Price

Chapter 1. GIS Data

Objectives

- Understanding how real-world features are represented by GIS data
- Knowing the differences between the raster and vector data models
- Getting familiar with the basic elements of data quality and metadata
- Learning the different types of GIS files used by ArcGIS
- Using ArcMap to view GIS data
- Learning about layers and their properties

Mastering the Concepts

GIS Concepts

Representing real-world objects on maps

To work with maps on a computer requires developing methods to store different types of map data and the information associated with them. Objects in the real world, such as cities, roads, soils, rivers, and topography, must first be portrayed as map objects, such as those on a paper topographic map. These map objects must then be encoded for storage on a computer.

Many different data formats have been invented to encode data for use with GIS programs; however, most follow one of two basic approaches: the **vector** model or the **raster** model. In either approach, the critical task includes representing the information at a point, or over a region in space, using *x* and *y* coordinate values (and sometimes *z* for height). The *x* and *y* coordinates are the spatial data. The information being represented, such as a soil type or a chemical analysis of a well, is called the attribute data. Raster and vector data models both store spatial and attribute data, but they do it in different ways.

Both data systems are **georeferenced**, meaning that the information is tied to a specific location on the earth's surface using *x-y* coordinates defined in a standard way: a **coordinate system**. One can choose from a variety of coordinate systems, as we will see in Chapter 11. As long as the coordinate systems match, we can display any two spatial data sets together and have them appear in the correct spatial relationship to one another.

The vector model

Vector data use a series of *x-y* locations to store information (Fig. 1.1). Three basic vector objects exist: points, lines, and polygons. These objects are called **features**. **Point** features are used to represent objects that have no dimensions, such as a well or a sampling locality. **Line** features represent objects in one dimension, such as a road or a utility line. **Polygons** are used to represent two-dimensional areas, such as a parcel or a state. In all cases, the features are represented using one or more

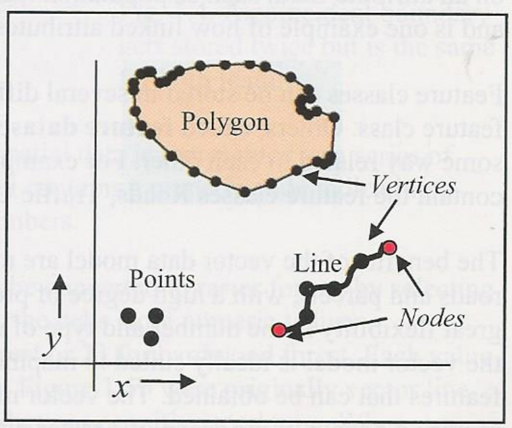

Fig. 1.1. The vector data model uses a series of *x-y* locations to represent points, lines, and polygon areas.

x-y coordinate locations (Fig. 1.1). A point consists of a single *x-y* coordinate pair. A line includes two or more pairs of coordinates—the endpoints of the line are termed **nodes**, and each of the intermediate points is called a **vertex**. A polygon is a group of **vertices** that define a closed area.

The type of object used to represent features depends on the scale of the map. A river would be represented as a line on a map of the United States because at that scale it is too small for its width to encompass any significant area on the map. If one is viewing a USGS topographic map, however, the river encompasses an area and would be represented as a polygon.

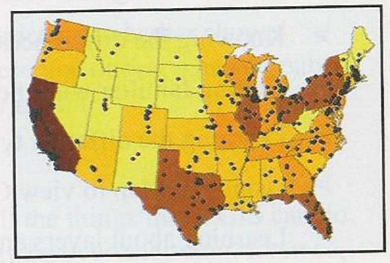

Fig. 1.2. A states feature class and a cities feature class

In GIS, like features are grouped into data sets called **feature classes** (Fig. 1.2). Roads and rivers are different types of features and would be stored in separate feature classes. A feature class can contain only one kind of geometry—it can include point features, line features, or polygon features but never a combination. In addition, objects in a feature class have information stored about them, such as their names or populations. This information is called the **attributes** and is stored in a table (Fig. 1.3). A special field, called the Feature ID (**FID**) or ObjectID (**OID**), links the spatial data with the attributes. Each feature's attributes are stored in one row of the table, and each column is a different type of information, such as population or area. A river and a highway would not be found in the same feature class because their information would be different—flow measurements for one versus pavement type for the other—and would need to be stored in different tables with different columns.

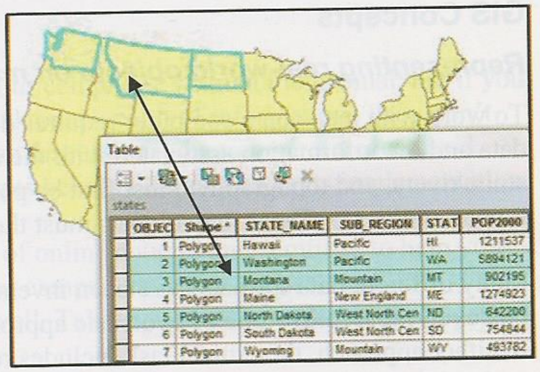

OBJEC	Shape *	STATE_NAME	SUB_REGION	STAT	POP2000
1	Polygon	Hawaii	Pacific	HI	1211537
2	Polygon	Washington	Pacific	WA	5894121
3	Polygon	Montana	Mountain	MT	902195
4	Polygon	Maine	New England	ME	1274923
5	Polygon	North Dakota	West North Cen	ND	642200
6	Polygon	South Dakota	West North Cen	SD	754844
7	Polygon	Wyoming	Mountain	WY	493782

Fig. 1.3. Each state is represented by a spatial feature (polygon), which is linked to the attributes.

When a state is highlighted on the map, its matching attributes are highlighted in the table, and vice versa. It is this live link between the spatial and attribute information that gives the GIS system its power. It enables us, for example, to create a map in which the states are colored based on an attribute field, such as population (Fig. 1.2). This technique is called **thematic mapping** and is one example of how linked attributes can be used to analyze geographic information.

Feature classes can be stored in several different formats. Some data formats only contain one feature class. Others, called **feature datasets**, can contain multiple feature classes that are in some way related to each other. For example, a feature dataset called Transportation might contain the feature classes Roads, Traffic Lights, Railroads, and Canals.

The benefits of the vector data model are many. First, it can store individual features, such as roads and parcels, with a high degree of precision. Second, the linked attribute table provides great flexibility in the number and type of attributes that can be stored about each feature. Third, the vector model is ideally suited to mapmaking because of the high precision and detail of features that can be obtained. The vector model is also a compact way of storing data, typically requiring a tenth of the space of a raster with similar information. Finally, the vector model is ideally suited to certain types of analysis problems, such as determining perimeters and areas, detecting whether features overlap, and modeling flow through networks.

However, the vector model has some drawbacks. First, it is poorly adapted to storing continuously varying surfaces, such as elevation or precipitation. Contour lines (as on topographic maps) have been used for many years to display surfaces, but calculating derived information from contours, such as slope, flow direction, and aspect, is difficult. Finally, some types of analysis are more time-consuming to perform with vectors.

Modeling feature behavior with topology

Two basic vector models exist: **spaghetti models** and **topological models**. A spaghetti model stores features of the file as independent objects, unrelated to each other. Simple and straightforward, this type of model is found in many types of applications that store spatial data. It is also commonly used to transfer vector features from one GIS system to another.

A topological data model stores features, but it also contains information about how the features are spatially related to each other. Many types of spatial relationships might be of interest, for example, whether two parcels share a common boundary (**adjacency**), whether two water lines are attached to each other (**connectivity**), whether a company sprayed pesticide over the same area on two different occasions (**overlap**), or whether a highway connects to a crossroad or has an overpass (**intersection**). Although computer algorithms can determine whether these spatial relationships exist between features in a spaghetti model, storing explicit information about the relationships can save time if the relationships must be used repeatedly.

Another application of topology involves analyzing the **logical consistency** of features. Logical consistency evaluates whether a data model or data set accurately represents the real-world relationships between features. For example, two adjacent states must share a common boundary that is exactly the same (the real-world situation), even though the states are stored in the data model as two separate features with two boundaries that coincide (Fig. 1.4). Lines representing streets should connect if the roads they represent meet. A line or a polygon boundary should not cross over itself.

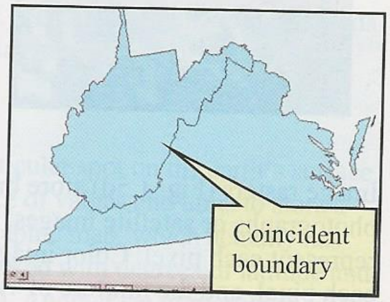

Fig. 1.4. A coincident boundary gets stored twice but is the same for both features.

Finally, topology can be used to better model the real-world behavior of features. In a network topology, for example, the connections between features are tracked so that flow through the network can be analyzed. Applications of networks include water in streams, traffic along roads, flights in and out of airline hubs, and utilities through pipes or electrical systems.

The raster model

The raster model has the benefit of simplicity. A set of spatial data is represented as a series of small squares, called **cells** or **pixels** (Fig. 1.5). Each pixel contains a numeric code indicating a single attribute, and the raster is stored as an array of numbers.

Vector features, such as roads or land use polygons, can be converted to raster format by selecting a single attribute to be stored in the cells. In Figure 1.5a, the cells store numeric values representing a land cover type, such as 46 for conifer forest or 23 for hardwood forest. Each value is given a different color for display. The roads shown in Figure 1.5b were originally vector line features with a text attribute indicating a primary, secondary, or primitive road type. When converted to a raster, the number (1, 2, or 3) is used to represent each road that passes through a cell. The cells that don't contain roads are given a null value. Rasters that store vector features in a raster format are sometimes called **discrete** rasters, because they represent discrete objects (roads, pipelines, land use polygons).

However, rasters may also be used to store a map quantity rather than features. Map quantities are values or variables that change over the earth's surface. A **digital elevation model** (DEM), for example, stores elevation values (Fig. 1.5c). Cells are unlikely to have the same elevation as their neighbors, and the values range smoothly into one another, forming a continuous surface (or continuous field). Therefore, they are commonly called **continuous** rasters.

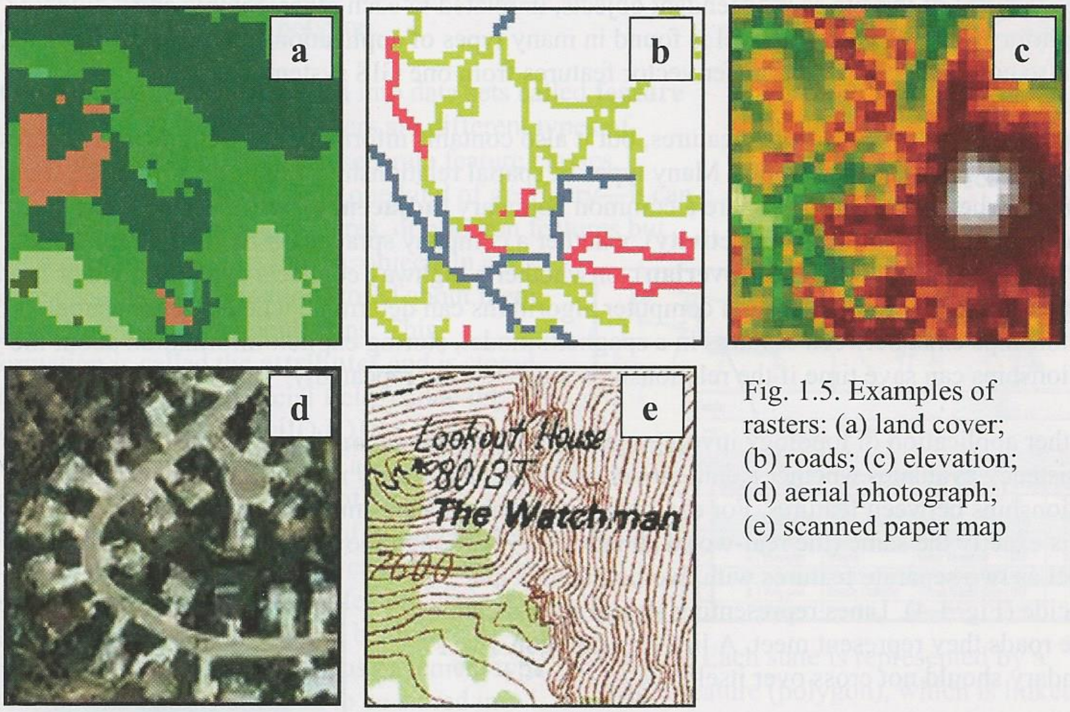

Fig. 1.5. Examples of rasters: (a) land cover; (b) roads; (c) elevation; (d) aerial photograph; (e) scanned paper map

Image rasters (Fig. 1.5d) store brightness values and are commonly used to store aerial photographs or satellite images. Images may contain multiple arrays of values, called **bands**, to represent each pixel. Color images often contain a red band, a green band, and a blue band for each pixel, and the mixture of values in each band defines the color of the pixel. One can also do a digital scan of a paper map and store it as a raster, such as the US Geological Survey topographic map shown in Figure 1.5e, known as a **digital raster graphic** (**DRG**). Each cell stores an index code representing a different color, such as 5 for the brown contours and 1 for the white background.

A raster data set is laid out as a series of rows and columns. Each pixel has an "address" indicated by its position in the array, such as row = 3 and column = 6. Georeferencing a map in an *x-y* coordinate system requires four numbers: an *x-y* location for one pixel in the raster data set and the size of the pixel in the *x* and *y* directions. Usually the upper-left corner is chosen as the known location, and the *x* and *y* pixel dimensions are the same so that the pixels are square. From these four numbers, it is possible to calculate the coordinates of every other pixel based on its row and column position. In this sense, the georeferencing of the pixels in a raster data set is implicit—one need not store the *x-y* location of every pixel.

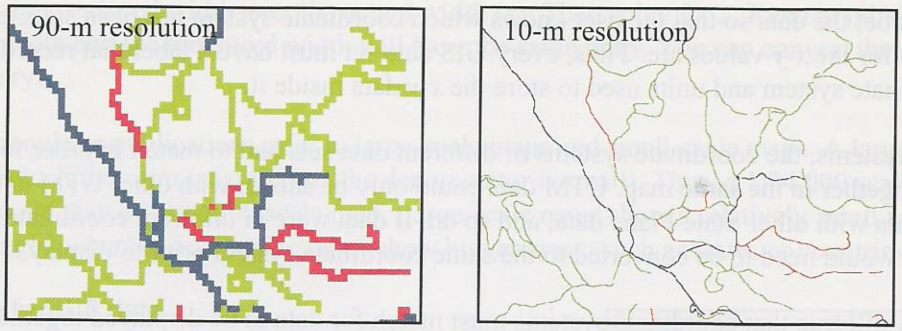

Fig. 1.6. Impact of raster resolution. The 10-meter resolution raster stores the roads more precisely, but it takes 81 times as much space.

The *x* and *y* dimensions of each pixel define the **resolution** of the raster data. The higher the resolution, the more precisely the data can be represented. Consider the 90-meter resolution roads raster in Figure 1.6. Since the raster cell dimensions are 90 meters, the roads are represented as much wider than they actually are, and they appear blocky. A 10-meter resolution raster could represent the roads more accurately; however, the file size would increase by 9×9, or 81 times.

The raster model mitigates some of the drawbacks of vectors. It is ideally suited to storing continuous and rapidly changing discontinuous information because each cell can have a value completely different from its neighbors. Many analyses are simple and rapid to perform, and an extensive set of analysis tools for rasters far outstrips those available for vectors.

The drawbacks of rasters lie chiefly in two areas. First, they suffer from trade-offs between precision and storage space to a greater extent than vectors do. The second major drawback of rasters is that they can store only one numeric attribute per raster. Vector files can store hundreds of attribute values for each spatial feature and can handle text data more efficiently.

Coordinate systems

Both raster and vector data rely on *x-y* values to locate data to a particular spot on the earth's surface. The *x-y* values of the coordinate pairs can vary, however. The choice of values and units to store a data set is called its **coordinate system**. Consider a standard topographic map, which actually has three different coordinate systems marked on it. The corners are marked with degrees of latitude and longitude, which is termed a **geographic coordinate system (GCS)**. Another set of markings indicates a scale in meters representing the UTM, or Universal Transverse Mercator, coordinate system. A third set of markings shows a scale in feet, corresponding to a State Plane coordinate system. Any location on the map can be represented by three different *x-y* pairs corresponding to one of the three coordinate systems (Fig. 1.7). A global positioning system (GPS) unit also has this flexibility. It can be set to record a location in GCS degrees, UTM meters, State Plane feet, or other coordinate systems.

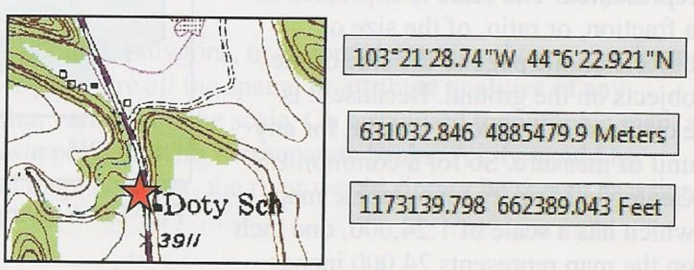

Fig. 1.7. A location can be stored using different coordinate systems and units. The *x-y* location of this school is shown in degrees, UTM meters, and State Plane feet.

When creating a vector or raster data set, one must choose a coordinate system and units for storing the *x-y* values. It is also important to label the data so that the user knows which coordinate system has been selected and what the units for the *x-y* values are. Thus, every GIS data set must have a label that records the type of coordinate system and units used to store the *x-y* data inside it.

In older GIS systems, the coordinate systems of different data sets had to match in order for them to be drawn together in the same map. UTM data could only be shown with other UTM data, State Plane data with other State Plane data, and so on. If data were in different coordinate systems, they would need to be converted to the same coordinate system prior to display.

Although it is still true that coordinate systems must match for data to be displayed together, many GIS systems can now perform the conversion on-the-fly. This feature allows data to be stored in different coordinate systems yet to be drawn together. In ArcMap, the user defines a coordinate system for the map, and all of the data are converted to match (Fig. 1.8). The units defined for the map coordinate system, whether they are meters, feet, or degrees, become the **map units** and may differ from the stored units in the files. In Figure 1.8, the UTM data use meters to store the *x-y* coordinates, the GCS data use degrees, and the State Plane data use feet. The Oregon Statewide Lambert coordinate system uses meters, so meters become the map units.

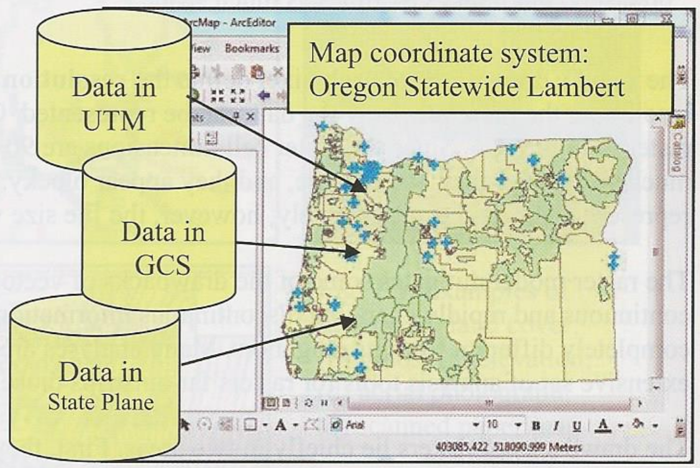

Fig. 1.8. Data in any coordinate system can be displayed together by setting the data frame coordinate system.

Map scale

The act of taking a set of GIS features with *x-y* coordinate values and drawing them on a screen or printing them on a piece of paper establishes a map scale. On a paper map, the scale is fixed at the time of printing. Within a computer system that allows interactive display, the scale changes every time the user zooms in or out of the map.

What is map scale?

Map scale is a measure of the size at which features in a map are represented. The scale is expressed as a fraction, or ratio, of the size of objects on the page to the size of the objects on the ground. Because it is expressed as a ratio, it is valid for any unit of measure. So for a common US Geological Survey topographic map, which has a scale of 1:24,000, one inch on the map represents 24,000 inches on the ground. Imagine a map made to the scale of 1:100,000. You can use the map scale and a ruler to determine the true distance of any feature on the map, such as the width of a lake (Fig. 1.9). Measure the lake with a ruler and then set up a proportion such that the map scale

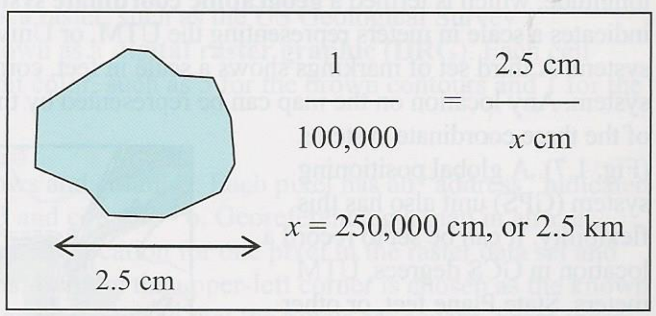

Fig. 1.9. Solving for the size of a lake.

equals the measured width over the actual width (x). Then solve for x. Keep in mind that the actual width and the measured width will have the same units. You can convert these units, if necessary.

Often people or publications refer to large-scale maps and small-scale maps. A large-scale map is one in which the *ratio* is large (i.e., the denominator is small). Thus, a 1:24,000 scale map is larger scale than a 1:100,000 scale map. Large-scale maps show a relatively small area, such as a quadrangle, whereas small-scale maps show bigger areas, such as states or countries.

Scales for GIS data

When data are stored in a GIS, they technically do not have a scale because only the coordinates are stored. They acquire a scale once they are drawn on the screen or on a piece of paper. However, most data layers have an intrinsic scale at which they were created. A 1:1 million scale paper map that is scanned or digitized cannot effectively be used at larger scales. The map in Figure 1.10 shows congressional districts in pink and the state outlines in thick black lines. The state boundaries are more angular and less detailed than the districts because they were digitized at a smaller scale. Thus, although it is possible to take small-scale data and zoom in to large scales, the accuracy and detail of the data will suffer. The original scale of a data set is an important attribute and is included when documenting it.

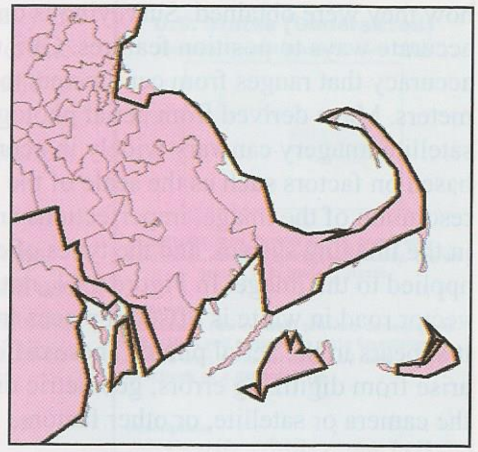

Fig. 1.10. These two layers showing Massachusetts originated from maps with different scales.

One should exercise caution in using data at scales very different from the original source. Zooming in to a data set may give a false impression that the data are more precise than they actually are. A pipeline digitized from a 1:100,000 scale map has an uncertainty of about 170 feet in its location due to the thickness of the line on the paper. Displaying the pipeline on a city map at 1:60,000 might look fine, but zooming to 1:2000 would not help, should you desire to locate the pipeline by digging.

From looking at Figure 1.10, one might conclude that it is desirable to always obtain and use data at the largest possible scale. However, large-scale data require more data points per unit area, increasing data storage space and slowing the drawing of layers. Every application has an optimal scale, and little is gained by using information at a higher scale than needed.

Data quality

Representing real-world objects as points, lines, polygons, or rasters always involves some degree of generalization. No data file can exactly capture all the spatial or attribute qualities of any object. The degree of generalization often varies with the scale. On a standard topographic map, a river has a width and can be modeled as a polygon with two separate banks. A city would be shown as a polygon area. For a national map, however, the river would simply be shown as a line, and a city would be shown as a point.

Even a detailed representation of an object is not always "true." Rivers and lakes can enlarge in size during a spring flood event or shrink during a drought. The boundary of a city changes over time as the city grows. Users of GIS data must never forget that the data they collect and use will contain flaws, and that the user has an ethical and legal responsibility to ensure that the data used for a particular purpose are appropriate for the task. When evaluating the quality of a data set, geospatial professionals consider the following aspects.

Geometric accuracy refers to the *x-y* values of a feature class or raster. How closely do the locations correspond to the actual location on the earth's surface? Geometric accuracy is usually a function of the original scale at which data are collected and of how they were obtained. Surveying is one of the most accurate ways to position features. GPS units have an accuracy that ranges from centimeters to tens of meters. Maps derived from aerial photography or satellite imagery can vary widely in geometric quality based on factors such as the scale of the image, the resolution of the image, imperfections and distortions in the imaging system, and the types of corrections applied to the image. In Figure 1.11, notice that the vector road in white is offset in places from the road as it appears in the aerial photo. These differences can arise from digitizing errors, geometric distortions from the camera or satellite, or other factors.

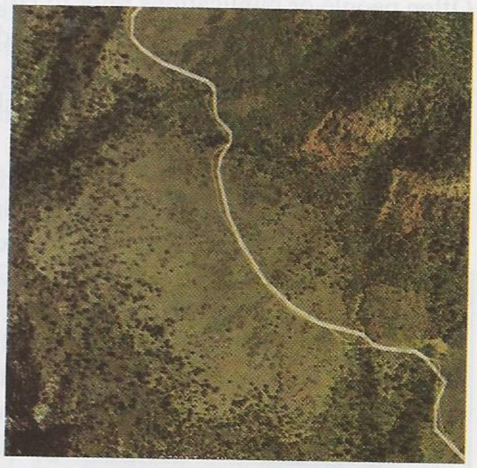

Fig. 1.11. Aerial photo near Woodenshoe Canyon, Utah. Source: Google Earth and TeleAtlas.

Moreover, not every boundary can be as precisely located as a road. Imagine that you wish to delineate the land-cover types: *forest*, *shrubland*, *grassland*, and *bare rock* in this photo. Where would you draw the line between *shrub* and *grassland*? At what point does the *shrubland* become *forest*? Six people given this photo would come up with six different maps. Some boundaries would match closely; others would vary as each person made a subjective decision about where to place each boundary.

Thematic accuracy refers to the attributes. Some types of data are relatively straightforward to record, such as the name of a city or the number of lanes in a road. Even in this situation, the value of a feature might be incorrectly recorded. Other types of information can never be known exactly. Population data, for example, are collected through a process of surveying and self-reporting that takes many months. It is impossible to include every person. Moreover, people are born and die during the survey process, or are moving in and out of towns. Population data can never be more than an estimate. These difficulties don't mean that it is pointless to collect the data. However, it is important to understand the limitations and potential biases associated with thematic data.

Resolution refers to the sampling interval at which data are acquired. Resolution may be spatial, thematic, or temporal. Spatial resolution indicates at what distance interval measurements are taken or recorded. What is the size of a single pixel of satellite data? If one is collecting GPS points by driving along a road, at what interval is each point collected? Thematic resolution can be impacted by using categories rather than measured quantities: if one is collecting information on the percent crown cover in a forest, is each measurement reported as a continuous value (32%) or as a classified range (low, medium, high)? Temporal resolution indicates how frequently measurements are taken. Census data are collected every 10 years. Temperature data taken at a climate station might be recorded every 15 minutes, but it might also be reported as a monthly or yearly average.

Precision refers to either the number of significant digits used to record a measurement or the statistical variation of a repeated single measurement. Many people confuse precision with accuracy, but it is important to understand the distinction. Imagine recording your body temperature with an oral digital thermometer that records to a thousandth of a degree and getting the value of 99.894 degrees Fahrenheit. This measurement would be considered precise. However, imagine that you take the reading immediately after drinking a cup of hot coffee. This action throws off the thermometer reading so that it does not record your true body temperature. Thus, the measurement is precise, but it is not accurate.

Evaluating the quality of a data set can be difficult, especially if the data were created by someone else. Professionals who create data incur an obligation to evaluate the quality of the data and to provide a report that summarizes the spatial and thematic accuracy so that users can determine whether a data set is suited to a particular purpose. Producers should also provide information on other aspects of a data set, such as what geographic area it covers, what coordinate system it uses, what the information in the attribute tables means, how a potential user can access the data, and more. If the original data were created or compiled by others, the producer must also give proper credit to the originators. Such information about a data set is called **metadata** (Fig. 1.12). The content and format of metadata are established by international standards.

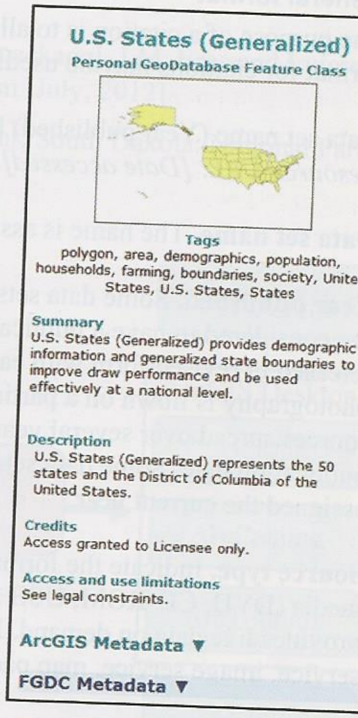

Fig. 1.12. Metadata

Organizations assemble collections of metadata to allow potential users to search for and evaluate data sets before they are downloaded or purchased, like an electronic library catalog. Once a candidate data set is identified, the user can explore the full metadata record to determine if the data set is appropriate for the particular application. If so, the metadata tell the user where the data are located, how they can be obtained, and what the cost might be.

Metadata also record the access and use constraints on GIS data. GIS data can be copyrighted and their uses restricted to certain people or certain actions. Some GIS data, including most data sets derived from federal agencies, can be freely copied and redistributed with credit given to the originating agency. Other data are developed by companies, and the rights are licensed to specific users. Often the license includes the right to distribute maps or other static copies derived from the data, but not the data themselves. Every user is responsible for understanding the applicable use constraints placed on any data set and for abiding by them. Failure to do so can result in civil and criminal penalties against the individual or the organization he or she works for.

Citing GIS data sources

Ethical and professional considerations require that any map, publication, or report should cite the data source(s) used and give proper credit to the originators of the data. The metadata, or the site from which the data were obtained, are good sources of information for citations. The best practice is to record the citation when the data are obtained so the information is available when needed. Generally one cites only data that are publically available (free or purchased). Data created internally within the workplace need not be cited, although often the company name or logo will appear on the map. A data set provided once in response to a personal request should be cited as a personal communication.

Keep in mind, when citing, that the place you found the data may not always be its source. Your GIS administrator may have placed often-used data sets, such as ESRI Data and Maps, on a workplace fileserver for easy access, but the fact that they were obtained internally does not free you from the need to cite them, and you must cite the original source, not the local server.

General format

The purpose of a citation is to allow people to obtain the data if they wish. The following general format for citations may be used:

Data set name (Year published) [source type]. Producer name, producer contact information. *Resource URL: [Date accessed].*

Data set name. The name is assigned by the creator or provider of the data.

Year published. Some data sets are assembled and provided once or at long intervals, and these are considered to have a publication date. For example, the ESRI Data and Maps product is released in revised form with each version of the software and carries a publication year. Aerial photography is flown on a particular date (although mosaics such as Google Earth use multiple sources spread over several years). It may take a little hunting or a few questions to find the publication date. Some data sets are updated at shorter intervals, or are even live. These should be assigned the current year.

Source type. Indicate the format in which the data are available. Types might include physical media (DVD, CD-ROM, USB stick), a file downloaded from the Internet, or a service that provides live data on demand. Different types of services exist, such as database service, map service, image service, map package, and layer package. New types are being added all the time.

Producer name. Give the name of the person or agency that makes the data available. In some cases, this may be different from the originator of the data. For example, ESRI publishes Data and Maps using data from many different sources. ArcGIS Online serves public data.

Producer contact information. For a company or small agency, the city and state should be included. For large agencies, particularly those with many offices but a unified web presence, the name itself is sufficient. Indicate a clearing house name, such as ArcGIS Online, if appropriate.

Resource URL. This entry is optional; include when appropriate. Use only static URLs even if it means the user has to hunt for the data. (A static URL always has the same form, whereas a dynamic URL is generated automatically based on search strings or other information. Dynamic URLs generally contain gibberish and/or characters like %.)

Date Accessed. This entry is optional because it mainly applies to online data sets. Include the year and month when you accessed or downloaded the data.

Examples of citations

Black Hills National Forest Database (2008) [downloaded file]. Black Hills National Forest, Custer, SD. URL: http://www.fs.usda.gov/main/blackhills/landmanagement/gis [August, 2010].

ESRI Data and Maps (2010) [DVD]. ESRI, Inc., Redlands, CA.

National Hydrology Dataset (2012) [downloaded file] United States Geological Survey on the National Map Viewer. URL: http://viewer.nationalmap.gov/viewer/ [July 23, 2012].

Geographic Names Information System (2008) [downloaded file]. United States Geological Survey. URL: http://geonames.usgs.gov/domestic/download_data.htm [May 21, 2009].

USA Topo Maps (2009) [map service]. ESRI on ArcGIS Online. URL: http://server.arcgisonline.com/ arcgis/services/USA_Topo_Maps/MapServer [January 1, 2012].

EIA Coalbed Methane Field Boundaries (2011) [map service]. US Department of Energy on ArcGIS Online. URL: http://arcgis.com [August, 2013].

Mineral Operations of Africa and the Middle East (2010) [layer package]. J.M. Eros and Luissette Candelario-Quintana on ArcGIS Online. URL: http://ArcGIS.com [July, 2012].

Badlands National Park GIS Database (2012) [CD-ROM]. Interior, South Dakota: National Park Service—Badlands National Park, personal communication.

About ArcGIS

ArcGIS overview

ArcGIS is developed and sold by Environmental Systems Research Institute, Inc. (ESRI). It has a long history and has been through many versions and changes. Originally developed for large mainframe computers, in the last 15 years it has metamorphosed from a system based on typed commands to a user-friendly graphical user interface (GUI). Data models, too, have changed over time, so that one is likely to encounter data sets in different formats. Knowing this background helps a student of GIS understand the nature of the ArcGIS system and its data.

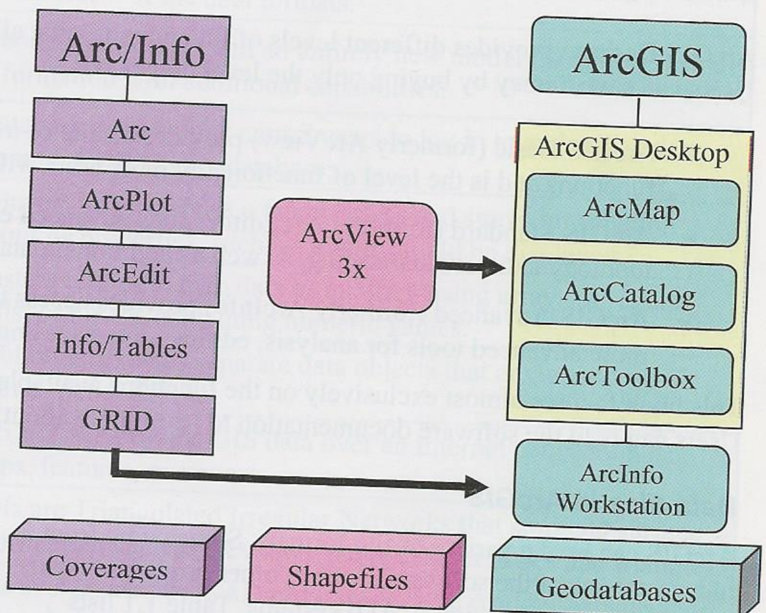

Fig. 1.13. Relationship between ESRI products and data formats

The old core of the ArcGIS system was called Arc/Info and included three programs: Arc, ArcEdit, and ArcPlot (Fig. 1.13), which utilized a data model known as a **coverage** and incorporated a database program called INFO, which appears primitive today. All of the programs were command based, meaning that the user typed commands into a window to make the program work.

The difficulty of learning Arc/Info prompted ESRI to create ArcView in 1992, which was easier to use but not as powerful as Arc/Info. ArcView was designed primarily to view and analyze spatial data rather than create them. ArcView used a simpler data model, called the shapefile, although it could read coverages and convert them to shapefiles.

ArcGIS Desktop, released in 2001, combines power and ease of use. It contains two main programs.

- ArcMap provides the means to display, analyze, and edit spatial data and data tables.
- ArcCatalog facilitates viewing and managing of spatial data files. It should *always* be used to delete, copy, rename, or move spatial data files.

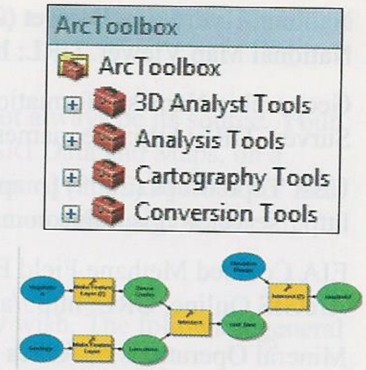

Fig. 1.14. ArcToolbox and a model

In addition, ArcGIS Desktop contains ArcToolbox, a collection of tools and functions for operations in ArcCatalog and ArcMap, such as converting between data formats, managing map projections, and performing analysis (Fig. 1.14). Users may create and add their own tools or scripts for special or often-used tasks. A program called Model Builder lets users graphically arrange and run sequences of steps, and save them to be used over and over.

ArcGIS Desktop provides different levels of functionality that all use the same basic interface. Users can save money by buying only the level they need.

- ArcGIS Basic (formerly ArcView) provides a suite of mapping, editing, and analysis functions and is the level of functionality most users will require on a regular basis.
- ArcGIS Standard (formerly ArcEditor) adds advanced editing capabilities, such as topology and network editing, as well as additional data conversion tools.
- ArcGIS Advanced (formerly ArcInfo) provides access to the full functionality and adds more advanced tools for analysis, editing, data handling, and cartography.

This book focuses almost exclusively on the functions available with an ArcGIS Basic license. Users can read the software documentation to learn more about the advanced capabilities.

Data files in ArcGIS

ArcGIS can read a variety of file formats. Some come from older versions of the software, or from other programs such as computer-aided design (CAD) systems. Table 1.1 lists many of the data sets that can be used in ArcGIS.

Shapefiles

Shapefiles are spaghetti data models developed for the early version of ArcView. A shapefile contains one feature class composed of points or lines or polygons but never a mixture. The attributes are stored in a **dBase** file. Shapefiles can, however, store **multipart features**, which are single features made of multiple objects. For example, Hawaii requires multiple polygons to represent the different islands, but the state can be stored as a multipart feature so that it has only one record in the attribute table.

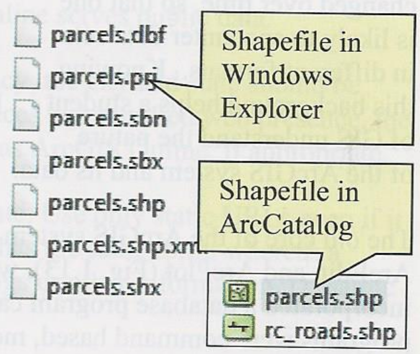

Fig. 1.15. Shapefiles are groups of files but appear as single entries in ArcCatalog.

Although a shapefile appears as one icon in ArcGIS, it is actually composed of multiple data files, as can be seen in Windows Explorer (Fig. 1.15). The rivers shapefile consists of seven different files. The .shp file stores coordinate data, the .dbf file stores attribute data, and the .shx file stores a spatial index that speeds drawing and analysis. These first three files are required for every shapefile to function. Additional files may also be present: the .prj file stores projection information, the .avl file is a stored legend, and the .xml file contains metadata. To copy a shapefile, all of these files must be included. ArcCatalog takes care of this automatically, but Windows Explorer does not—one reason it is wise to manage GIS data using ArcCatalog.

In a shapefile attribute table, the first two columns of data are reserved for storing the feature identification code (FID) and the coordinate geometry (Shape) field. These fields are created and maintained by ArcGIS and cannot be modified by the user. All other fields are added by the user.

Table 1.1. Types of files and data sources used by ArcGIS

File type	Description
Shapefiles	Shapefiles are vector feature classes developed for the early version of ArcView and have been carried over into ArcGIS.
Coverages	A coverage is the vector data format developed for Arc/Info and is the oldest of the data formats.
Geodatabases	Geodatabases represent an entirely new model for storing spatial information with additional capabilities.
Database connections	Database connections permit users to log in to and utilize data from an RDBMS geodatabase.
Layer files	A layer file references a feature class and stores information about its properties, such as how it should be displayed.
Rasters	Rasters represent map data or imagery using arrays of regular cells, or pixels, containing numeric values.
Tables	Tables can exist as separate data objects that are unassociated with a spatial data set.
GIS servers	GIS Servers provide GIS data over an Internet connection as maps, features, or images
TINs	TINs are Triangulated Irregular Networks that store 3D surface information, such as elevation, using a set of nodes and triangles.
CAD drawings	Data sets created by CAD programs can be read by ArcGIS, although they cannot be edited or analyzed unless they are converted to shapefiles or geodatabases.

Geodatabases

A geodatabase can contain many different objects, including multiple feature classes, geometric networks, tables, rasters, and other objects. Figure 1.16 shows a geodatabase named rapidnets. Feature classes may exist as individual objects in a geodatabase (as do the restaurants or schools), or they may be grouped into feature datasets. A feature dataset contains a collection of related feature classes with the same coordinate system, such as the Utilities feature dataset in Figure 1.16.

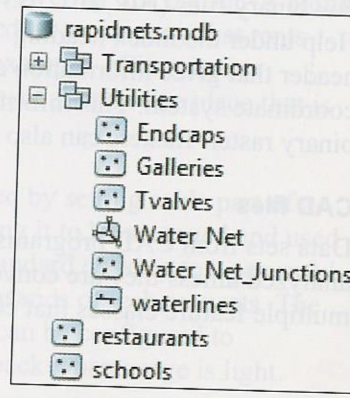

Fig. 1.16. A geodatabase

A feature dataset can also store topological associations between feature classes. The Utilities feature dataset in Figure 1.16

contains a **network topology** stored in the Water_Net and Water_Net_Junctions feature classes. Chapter 9 covers some special analysis functions that can be used with networks. Feature datasets may also contain **planar topology**, which contains rules about spatial relationships within or between layers. Special tools help find and fix errors that violate the topology rules.

Finally, geodatabases may contain rules that assist in entering and validating attribute data. Called **domains**, these rules specify which values or range of values may be entered in a particular field; a percent field, for example, should contain only numbers between 0 and 100.

Three types of geodatabases are used by ArcGIS: personal geodatabases, file geodatabases, and SDE geodatabases. The behavior of the three types is similar, but the data storage formats and capabilities differ. They are described in Chapter 14.

Coverages

A **coverage** is the oldest vector format, developed for Arc/Info. ArcGIS Desktop has limited functions for managing coverages, so most users will encounter them simply as an old data format that must be converted to a shapefile or exported to a geodatabase. Several things are helpful to know in this process.

Coverages contain multiple feature classes, which may store points, arcs, polygons, and polygon labels. Coverages also store topology, and the tables have several attribute fields reserved for this purpose. Figure 1.17 shows these fields for a coverage called LANDUSE (some fields use the coverage name as part of the field name). It makes sense to delete these fields during or after the conversion, for they serve no useful purpose afterwards.

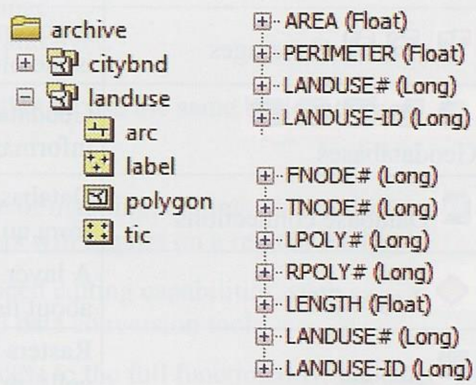

Fig. 1.17. Coverages contain fields of little value once they are converted.

VERY IMPORTANT TIP: Do not use Windows to copy or delete coverages, shapefiles, and geodatabases. Always use ArcCatalog to delete or copy spatial data sets.

Rasters

Rasters in ArcGIS can take a variety of formats, such as JPEG, TIF, GEOTIF, BMP, MrSID, and raw binary files (BIP, BIL, BSQ). A list of supported raster formats can be found in the ArcMap Help under the index heading "rasters, formats." Most rasters consist of the spatial data plus a header that gives information about the file, such as its number of rows and columns and its coordinate system. This information may be stored in a separate file or as the first part of the binary raster. Rasters can also be stored inside geodatabases.

CAD files

Data sets from CAD programs can be read by ArcGIS, although they cannot be edited or analyzed unless they are converted to shapefiles or geodatabases. A CAD file may contain multiple feature classes that correspond to layers of the drawing and can be opened separately.

Database connections

A user can connect to a database management system (DBMS) on a network through a database connection. This connection can be open, or it can require a login and password. Once inside, the user can access the tables within the database. Database connections may be available on a local network through a mounted drive, but some can be accessed online as database services.

Internet data services

Many organizations make data available over the Internet. Free data services, such as GoogleEarth and MapQuest, provide access to huge volumes of image and map data. These services are used online in a web interface and cannot be viewed in ArcMap. Although the data quality and documentation procedures are not designed for professional-level work, the volume and popularity of these sites introduce many people to GIS techniques and data. Other sites, such as the United States Geological Survey's National Map, do allow users to download data.

GIS servers

GIS servers provide geospatial data over Internet connections. Hosting GIS data on a server requires software that responds to requests from users for specific maps or data in the user's current window. Some GIS services are open and free, designed to provide public access to data and maps. Others may be locked down to a particular organization or group of users. Many organizations use servers to make in-house data available to personnel out in the field, or to allow employees in different locations to access the same company data sets. Servers provide data in a client-neutral format; as long as the client program knows how to use the service, the data can be used in ArcGIS Desktop, in a web application, on a tablet, or on a smartphone.

Several types of services may be offered. A **map service** renders map layers as tiles and sends them to the user as static images. This type of service is efficient and fast, but the user cannot modify how the map is symbolized. A **feature service** shares the requested data features; it tends to be slower, but the user can change how the features look and potentially edit or analyze them. A **layer package** or **map package** sends the features to the client, where they are stored locally during use. An **image service** provides access to large mosaics of satellite imagery or aerial photography. GIS servers can also provide analysis tools in addition to data. A **geoprocessing service** makes available certain computations and functions, so that users can perform the designated analysis through a web site, even if they do not have GIS software installed.

Cloud-based services

Setting up and managing a GIS server is a complex and expensive task requiring suitable space, equipment, software, and expertise. New cloud-based services are lowering this barrier. A **cloud** consists of warehouses of giant computers and hard drives managed by a company that rents processing power and disk space to clients. Some companies, for example, offer space to individuals to back up their computers or store their movies and music files in one place that is accessible to all of their devices.

Another type of cloud service is the virtual machine, or VM, created by setting aside part of a large computer with its own operating system and software, allowing it to be managed and used as an ordinary physical computer. The VM can be scaled from a standard desktop to a powerful machine capable of running intensive computations or serving thousands of web requests. The client rents the VM for an hourly fee. Furthermore, cloud services can be configured to automatically add more VM units when demand is high and scale back when usage is light.

Many organizations have moved to the cloud to host their GIS data, either as a complement to or even instead of housing their own GIS servers. Advantages of cloud services include the speed and ease of deployment, the ability to scale up and down to meet changing demand, and the security benefits of hosting public information outside the organizational firewall.

ArcGIS Online and web maps

ArcGIS Online is a cloud-based platform that provides an environment to create and share maps. It makes data available to people with little to no GIS training, yet it also addresses the needs of professional users. Anyone may use the data, and those willing to set up a free account may also share their data. The user controls whether a published data set is visible to a few selected people or to the general public. The platform also contains easy tools for creating web pages that include interactive maps. Organizations may purchase a subscription account, which allows them to manage security, designate users allowed to access or publish data, and host data in the cloud without needing to configure their own GIS servers.

ArcGIS Online is designed around the **web map**, an interactive map based solely on GIS services. Web maps generally perform a restricted set of basic functions, such as zoom and query; however, they are device independent and can be used in ArcGIS Desktop, in web browsers, on mobile devices (such as smartphones or tablets), and even within social media sites. They are relatively simple to create and to share.

Summary

- A GIS is a database system that uses both spatial and attribute data to answer questions about where things are and how they are related. It has many functions, including creating data, making maps, and analyzing relationships.
- Raster data employ arrays of values representing conditions on the ground within a square called a pixel. The array is georeferenced to a ground location using a single *x-y* point.
- Vector data use sequences of *x-y* coordinates to store point, line, or polygon features. Every feature is linked to an attribute table containing information about the feature.
- Every GIS data set has a coordinate system defined for stored *x-y* coordinate values. Many different coordinate systems are used, so each data set must be labeled with information about the coordinate system.
- Data are stored as simple spaghetti models or as topological models. Topological models can better model feature behavior and aid in locating and correcting geometric errors.
- Every GIS user has a responsibility to ensure that data are suitable for the proposed application. Data quality is measured in terms of geometric accuracy, thematic accuracy, resolution, and precision.
- Metadata store information about GIS data layers to help people use them properly.
- ArcGIS Desktop employs a menu-based interface for two programs called ArcMap and ArcCatalog. It uses many data formats, including shapefiles, coverages, geodatabases, rasters, images, TINs, CAD drawings, and Internet-based data services.
- ArcToolbox contains functions for processing, managing, and analyzing GIS data. Users may customize it by building models or writing scripts to repeat often-used sequences.