

11.2 WHAT IS SPATIAL ANALYSIS?

You have read that various methods of geospatial analysis and modelling devised over the years. Many of them date before the advent of GIS, when calculations and measurements were performed manually. Numbers of methods have been provided by the developers of GIS software, specialised companies and individuals. Spatial analysis is the heart of GIS.

You have already been introduced to concepts of spatial data in Block 2 *Fundamentals of GIS* of MGY-003. The strength of the GIS lies in its analytical capabilities i.e., spatial analysis.

Spatial analysis deals with adding significance to the geographic data and converting the data into useful information. Spatial analysis allows us to study real-world processes including the present situation of specific areas and features or the change in situation or the trends e.g. decreasing or increasing urbanisation in Delhi and NCR (National Capital Region)? Thus the spatial analysis is a process that involves looking at geographic patterns in the data and their relationships. The actual methods for spatial analysis can be very simple such as just by making a map of the theme you are analysing or more complex, involving models that mimic the real world by combining many data layers.

Thus the spatial analysis relates to the use of the technological inputs in collecting, storing, retrieving, displaying, manipulating, managing and analysing the spatial information. It is an amalgamation of remote sensing, GPS, GIS, digital cartography and database management systems. You can consider spatial analysis as the collaboration between the computer and the human, in which both play vital roles. Thus spatial analysis essentially requires an intelligent user, not just a powerful computer. In all of these cases GIS is the engine that performs the necessary operations, under the guidance of the user issuing the necessary instructions. Spatial analysis can reveal things that might otherwise be invisible or hidden.

You have read and found that all the possible definitions of spatial analysis express the basic idea that information on locations is essential and also that the analysis carried out without knowledge of location is not spatial analysis. Spatial analysis is a set of methods whose results change when the locations of the objects being analysed change. For example analysis of the average savings for a group of people is in no way spatial analysis because it is not concerned with the location of the people.

Later in this unit you will learn about simplest capabilities of GIS: counting and locating objects. These sets of techniques resemble the descriptive work by early explorers. This old approach should not be discounted, as it provides the basis for later analytical techniques. However, we will discuss about replacing the devices like dots, grids, distances wheels and planimeters traditionally used for measuring the sizes, lengths and areas of objects with computer algorithm that perform the same techniques.

Let us say for example, you are interested to represent a road. It can be represented by recording an appropriate value in a swath of cells in a raster

representation. With a cell size of 1m, the swath corresponding to a major four-lane highway might be as much as 50 or even 100 cells wide (Fig. 11.1a). Alternatively, the road might be represented as a single line, or centerline (Fig. 11.1b) in a vector database. In this case its location would be recorded by specifying the coordinates of a series of points aligned along the road's center. Finally, the road might be represented as an area, encased by its edges (Fig. 11.1c).




Fig. 11.1: Three alternative representations of the same road, as: a) swath of raster pixels, b) a vector centerline, c) a vector area

However, spatial analysis is also the most misused subsystem of the GIS. Much of the abuse of the GIS analysis is the result of lack of understanding of the nature of the spatial data contained in the system.

11.3 SPATIAL ELEMENTS

Spatial objects identified in the real world are easily identifiable into four types: *Points*, *Lines*, *Areas*, and *Surfaces* (Fig. 11.2). Collectively they are known as *spatial elements*. You have already been introduced to the spatial elements in Block 3 of MGY-003. The spatial elements can represent most of the substantial natural and human phenomena that we come across on everyday basis. In general, points, lines, and areas are used to explicitly represent real-world objects, where as surfaces are mostly used for volumetric representation, such as to represent hills, valleys.

Point features are spatial phenomena each of which occurs at one location in space. Each feature is said to be discrete in that it can occupy only a given point in space at any time and considered to have *no spatial dimension* – no width or length. Example of such feature would be a house or a village.

Line features are conceptualised as occupying only a *single dimension* in coordinate space. Roads, rivers, are the examples of linear features.

Area features have *two dimensions* both length and width dimensions. Area is composed of series of lines that begin and end at the same location.

Surface, add the third dimension of height to area features. Surfaces have three dimensions – length, width, and height. For instance, hills, valleys, and ridges can

be described by citing their locations, amount of area they occupy, how they are oriented and by noting their heights.









Spatial Elements and Representation					
		Point representation	Line representation	Area representation	Surface representations
Real World Phenomena	Point objects	 Tree	 Boulders Boulder train	 Animals Animal range	 Housing density
	Line objects	 Airport	 Highway	 Stream Watershed	 Hedgerow density
	Area objects	 Chemical spill	 Right of way Power line	 New subdivision	 Acres undeveloped
	Volumetric objects	 Open-pit mine	 River valley River	 Irrigation drain	 Acre-feet of water

Fig. 11.2: Spatial Elements and their representation (DeMers, 2000)

11.4 VECTOR ANALYSIS

The operational models of the real world are built by making use of spatial datasets, based upon field and object. The vector data may be transformed into raster representation and vice versa. The transformation in most cases will result in the loss of information (e.g. resolution) and therefore such transformations may not be reversible. It is possible to separate methods of analysis of vector data sets from those that operate on raster data sets at a level. Since, there are very few instances of operations that require input of both kinds. There are so many possibilities, it is important to understand the ones that are used most often and to leave the less popular ones. Each method is described in terms of the problem it attempts to solve, the inputs required and the outputs that it generates.

11.4.1 Methods of Analysis for Points

You will begin here with the simplest kinds of geographic objects. Suppose you have records of each of the customers of an insurance agent, including each customer's location. Perhaps these originated in an address list and were subsequently converted to coordinates using the process known as *geocoding* or *address matching*. Dynamic segregation plots linearly referenced data on a coordinate system. Plotted on a map they might look something like Fig. 11.3. The two techniques are similar because both can locate data source that lacks x-, y- coordinates and both use linear features (e.g. streets, highways) as reference data. The data generated from geocoding and dynamic segregation is necessary for inputs for certain data analysis. Geocoded data are crucial for

crime mapping and analysis, and dynamic segregation layers are useful for managing and analysing highway related data.




Fig. 11.3: Locations of the customers of an insurance agent, shown in relation to the agent's office location and is ready for analysis

11.4.2 Measurement of Distance

Now in the next instance we might ask how far each customer is located from the agent's location. The agent might be interested in knowing how many customers do their business with the agent rather than with some other agent. What is the average distance between each of the customers and the agent and are there areas near the agent where advertising might increase the agent's market share? All of these questions require the ability to measure the distance between points, a commonly used function of a GIS. Distance measurement calculates distance between spatial features. To measure the distance between two points we need to know their locations in (x,y) coordinates. Coordinate systems were introduced to you in Unit 8, *Datums, Map Projections and Coordinate Systems* of MGY-001. You have read that the UTM (universal transverse Mercator) coordinate system is based on measurements of coordinates in meters, so distances calculated in this coordinate system are easily understood. In these cases the distance between two points (x_A, y_A) and (x_C, y_C) , representing the locations of the agent and a customer, are given by:

$$D = \sqrt{(x_A - x_C)^2 + (y_A - y_C)^2}$$

This equation is derived from Pythagoras's famous theorem, so it is also known as the Pythagorean distance or the length of a straight line between the two points. Pythagorean distance applies to points located on a plane not on the curved surface of the Earth. Therefore they should be used only when the points are close together. It is important to recognise that this problem of the Earth's curvature is not resolved by using latitude and longitude coordinates. Suppose we record location using latitude and longitude and plug these values into the Pythagorean equation as if they were y and x respectively. In this case the coordinates are measured in degrees, so the value of D will also be in degrees. There are many other bases for measuring distance because it is often necessary to allow for travel that must follow streets or avoid barriers of one kind or another.

11.4.3 Measurement of Area

The use of GIS is strongly supported with the argument that computers make it easier to take measurements from maps, as compared with manual methods of analysis of information shown on maps. Let us demonstrate this by the

measurement of an area. Suppose we want to measure the areas occupied by industry in a particular city and we are given a map showing the industrial locations in coloured areas. Manually, we have to use one of two traditional methods, known as dot counting and planimetry.

Dot counting proceeds by overlaying a transparent sheet covered by dots at a known *density* and counting the number of such dots falling within the area of interest. In order to get a reliable estimate, the high density of dots and smaller size is required. The process is tedious and may produce erroneous results.

Planimetry is a crude mechanical device to trace the outline of the area and the result is read off on the dial of the instrument. Again the process is tedious and error prone. Thus manual methods are frustrating and expensive in terms of the time required to obtain even a poor level of accuracy.

However this operation becomes simple if the measurement of area is represented in the digital form as polygons. A simple calculation is made for each straight edge of the polygon, using the coordinates of its endpoints and the calculations are summed around the polygon. Any vector GIS is able to do this quickly and accurately. The accuracy of the final result is limited only by the resolution of the original digitiser. The world's first GIS, the Canada Geographic Information System were developed in the mid 1960s precisely for this reason. It would have taken decades even after employing a large workforce, to produce the promised measurements manually. The results would have been of disappointing accuracy. Instead, a GIS was developed and all of the maps were digitised.

11.4.4 Buffer Operation

It might be interesting to outline on a map the area that is within 1km of the agent's location. Instead of ascertaining the distance of one customer from another in Fig. 11.3, you might in turn identify all of the points within a certain distance, taking a reference point. In GIS the term *buffer* is used and we say that the circle created by this operation constitutes the 1km buffer around the agent. Buffers can be created for any kind of object such as points, lines, or areas and are very widely used in GIS analysis. Fig. 11.4 shows buffers for each of these types of objects. *Buffering* creates buffer zone by measuring straight line distances from selected features.




Fig. 11.4: Buffers for (a) point, (b) line and (c) area objects

Buffer operation is important spatial analysis function in GIS. Let us read about the applications of buffering as mentioned below:

- A buffer zone is often treated as a protection zone and is used for planning and regulatory purposes.
- It is useful in locating the areas/population benefitted or denied of the facilities and essential services such as hospitals, post office, roads etc.
- It can also be used to undertake population and environmental studies. This kind of analysis is called *Proximity Analysis* wherein the buffer operation will generate polygon feature types.
- It helps in determining the spatial proximity or nearness of various geographic features.

11.4.5 Overlay

The major advantage of a GIS is that it allows to identify spatial relationships between map features and to create new relationships by integration of map layers on various themes. It is performed under the *Map Overlay*, which is a spatial analysis involving two or more maps. Each map can be thought of as a layer containing information on a particular theme. These maps can be integrated to produce an output map on the basis of some logical and / or arithmetic operation. An overlay operation combines the geometries and attributes of two feature layers to create the output. The geometry of the output represents the geometric intersection of the features from the input layers.

Overlay processes involves two (or more) layers. Merging is a simple overlay process that combines two or more layers into one. It leaves overlapping regions and does not create new attributed polygons where there is overlap. In contrast to a simple merge, each of these operations will produce a new layer with unique combinations of the input database polygons.

You will read here about the two overlay methods:

- Union
- Intersection

Union is the laying one GIS database on another to produce a combination of the two. The resulting GIS database will extend as far as both input GIS databases extend (Fig. 11.5). Thus union preserves all features from the inputs.

The area extent of the output combines the area extents of both input layers. Union requires that both input layers be polygon layers.




Fig. 11.5: Union overlay: The union operation is used in a vector overlay of two coverages: the first is a map of land ownership versus renting and the second is a landuse map showing agricultural versus urban. The result of this operation is a map of the entire study area, indicating agricultural and urban as well as owned versus rented land (source: modified from DeMers, 2000)

Intersection is the laying one GIS database on another to produce a combination of the two and finding the overlapping areas between two GIS databases (Fig. 11.6). Thus intersection preserves only those features that fall within the area extent common to the inputs.




Fig. 11.6: Intersection between two GIS databases shows common area produced from the combination

Intersection is the most preferred method of overlay because any feature on its output has the attribute data from both of its inputs.


Overlay operations can be classified into the following:

- 1) Points-in-Polygon
- 2) Line-in-Polygon
- 3) Polygon Overlay
- 1) **Points-in-Polygon**

A point-in-polygon operation is a key feature of GIS analysis. Point features of one layer can be overlaid on polygon features of another layer to identify the polygons within which each point falls. For example, if well sites are represented as points in a layer and land holdings as polygons in another layer. Point-in-Polygon overlay could be used to determine the well sites contained within each land holding.

The point-in-polygon operation is actually very simple to execute and you will read about this method briefly. One of the points in the Fig.11.7 will be used as

an example. A line is drawn from the point, in this case diagonally upwards. The number of intersections between this line and the boundaries of each polygon is counted. In the example, there is one intersection with the boundary of Polygon A, second with the boundary of Polygon B and third with the boundary of Polygon C. The polygon that contains the point is the only one with an odd number of boundary intersections (and it will always be true that exactly one polygon has an odd number - all other polygons will have an even number). One useful way to think about the result of the point in polygon operation is in terms of tables. The result of the point in polygon operation could be expressed in terms of an additional column for the census tract table shown in Fig. 11.7, recording the number of points found to lie in that polygon, as a preliminary to computing the local market penetration.



Census Tract	Number of Customers	Total households
1	2	249
2	8	377
3	7	401

Fig: 11.7: By using the point-in-polygon operation to identify the census tract containing each of the customer locations, it is possible to compare customer counts to other attributes of the tracts, such as the total number of households

2) Line-in-Polygon

In a line-in-polygon overlay operation, the output contains the same line features as in the input layer but each line feature is dissected by the polygon boundaries on an overlay layer. Part of line may also fall directly on a polygon boundary instead of within a polygon. Thus the output has more line segments than does the input layer. Each line segment on the output combines attributes from the input layer and the underlying polygon. For example, roads stored as line in one layer can be merged with country polygons from another layer using a Line-in-Polygon overlay. The result can be used to determine the number and type of roads that fall within the Country or State.

3) Polygon Overlay

Polygon Overlay is an important operation for vector GIS, but it has two distinct versions. You will read about the examples in this section, such that

each point on the map lies exactly in one polygon and overlay produces a similar map. Overlay in this version is conducted on an entire map or layer or coverage at a time.

Fig. 11.8 shows a typical application of polygon overlay. Let us take two input maps, e.g. a map of land use and a map of district boundaries. A reasonable question to ask of a GIS in a case like this is: "how much land in district A is in agricultural use?" To answer the question we must somehow compare the two maps, identifying the areas common both in district A and in agricultural land use and measure this area. You can do this using polygon overlay. The history of this method dates back to the 1960s, though efficient software to perform this operation on vector data sets was not developed until the late 1970s.

Fig. 11.8 shows the result of the polygon overlay operation, in the form of a new data set. Each of the polygons in the new data set represents a unique combination of the inputs and the boundaries. The resulting polygons are combinations of the lines from the two inputs.

Polygon overlay has a numerous uses, of which landuse is the simplest. Another popular use is population statistics.




Fig. 11.8: Polygon overlay, combining a map with three polygons and a map with two polygons to create a map with five polygons. The attributes of the new map combine the relevant attributes of the input maps

11.4.6 Reclassification

You know that before conducting any kind of study or survey it is necessary to classify them. Similar classification framework has long been extended for the analysis of Earth surface features by cartographers. Thousands of classification schemes produced has one thing in common; they all have a target audience or end user in mind. They enable you to explore your data, without making changes to the location of spatial elements and without creating new spatial elements. You can store data with GIS before they are classified and later on manipulate them within the GIS depending on your needs. A GIS provides a wide variety of ways to classify and reclassify the stored attribute data to achieve the result. The nature of working and applicability of the user can actually be called one of *reclassification*, because the data that are input to GIS have already been classified. Thus the GIS operator either displays the existing classification (a data storage and data retrieval operation) or manipulates the existing attributes to

create a classification more appropriate to the questions being asked and decisions to be made. Data retrieval involves the selective search of data.

Reclassification involves the (re)assignment of thematic values to categories of an existing map. *Reclassification* can be defined as repackaging of exiting information based on initial values, locations, size, shape, contiguity.

Let us read this with an example of data for household income from a town 'A', ranges from ₹ 0 to more than ₹ 75,000. Now you can reclassify this data into seven slabs and further into three major groups, as you can read below:

1.	0 – 5000	}	Low Income Group
2.	5000 – 10,000		
3.	10,000 – 20,000	}	Medium Income Group
4.	20,000 – 40,000		
5.	40,000 – 60,000		
6.	50,000 – 75,000	}	High Income Group
7.	> 75,000		

11.4.7 Query

The fundamental purpose of any Geographic Information System is to allow retrieval of desired information. This retrieval in GIS is achieved through 'Query' function. We will read about various types of queries mentioned below:

- **Spatial Query:** It includes queries made on the spatial properties of the geographic object. You can read about some spatial queries as mentioned below:

Query a location is represented by a Point. Locational query provides the record number, or point-ID number coordinates specific to spatial data model, latitude-longitude of the point and the projection coordinates.

Query Distance computes both the current and cumulative distances. There are several ways to provide the distance information, viz. the longitude and latitude of the interactive cursor location, the pixel resolution of the current window, great circle distance between two points etc.

Circular/Rectangular/Polygon Area Search are useful for querying a data layer and select only those entities that fall within a user specific radius / rectangular / polygon area.

- **Non-spatial (attribute) Query:** It refers to the characteristics of any feature or phenomenon which are stored in form of a table in a database with a unique identifier key. This 'attribute' table can be queried to locate row or column that satisfies the search conditions. The data linkage is established between the spatial elements and their corresponding attribute tables through the common feature identifier key.

Exact Matching is most often employed. Here the attribute table has a corresponding record entered for each of the spatial elements. Sometimes, few attribute information are collected in more detail than the other attributes; in that case *Non-exact Matching* is carried out. It is of two

types, *Hierarchical Matching* includes e.g., finance, unemployment data and population data; *Fuzzy Matching* comprises crop boundaries usually define the field edges that rarely match the boundary between the soil types

- **Topological query:** It is defined as the inter-relationship amongst the features that are dependent of distance or direction. The key topological relationships are connectivity, adjacency and containment.
- **Spatio-temporal query:** This query provides ready answers to simple questions of feature, areal histories and attributes for the past moments in time.
- **SQL based query formulation:** A Structured Query Language (SQL) is a tool to extract data from a database and present result in a useful manner. SQL is used extensively in many database applications and has become a standard for relational DBMS.


11.5 RASTER ANALYSIS

Spatial analysis consists of getting information using geospatial datasets. Here you will read about the spatial analysis with raster data. Raster GIS provides a very powerful basis for analysis, which is similar in many respects to the capabilities of other software that also relies on raster representations. For example, some of the raster operations described here will be familiar in concept to people who regularly use software for processing digital photographs or scanned documents, or for processing the images captured by remote sensing satellites. Vector data sets are similar in form to those used in computer-assisted design software (CAD) and in drawing software. In raster GIS the world is represented as a series of *layers*, each layer divides the same project area into the same set of rectangular or square cells or *pixels*.

This makes it very easy to compare layers, since the pixels in each layer exactly coincide. But it also means that the ability to represent detail is limited by a pixel size, since no feature on the Earth's surface is smaller than a single pixel that cannot be represented. Smaller features can be represented by reducing the pixel size, but only at the cost of rapidly increasing data volume. Each layer records the values of one variable or attribute, such as land use or county name, for each pixel. The recorded value might be the average value over the pixel, or the value at the exact centre of the pixel, or the commonest value found in the pixel.

If there are many attributes to record (the census, for example, reports hundreds of attributes for each county in India), then separate layers must be created for each attribute. However, some raster GIS allows more efficient solution to this particular problem, in which a single layer is used to record district ID for each pixel. Fig. 11.9 gives the many attributes corresponding to that ID. Raster operations include four basic types of operations:

- *local operations* processes the contents of data sets pixel by pixel, performing operations on each pixel or comparing the contents of the same pixel on each layer;
- *focal operations* compares the contents of a pixel with those of neighbouring pixels, using a fixed neighbourhood (often the pixel's eight immediate neighbours);
- *zonal operations* performs operations on zones or contiguous blocks of pixels having the same values;
- *global operations* are performed for all pixels. Examples of each type of operation are given in the following sections.



Districts	Population	Tax Rate
A	2019	6
B	4511	6.5%
C	7633	6.8%

Fig. 11.9: By using raster cells to store the IDs of polygons with multiple attributes in a related table it is possible to store many attributes in a single raster layer

You will read about them in details in the following sections. Often a raster data operation relates the input to the output with mathematical functions.

11.5.1 Local Operations

The local operations can compute the average of the input rasters on a cell by cell basis. The simplest kind of local operations occur when a single raster layer is processed. This is often done in order to apply a simple reclassification, such as: "All areas of Soil Classes 1, 3, and 5 are suitable building sites for residential development; all other areas are not." This could be operationalised by reclassifying the soil class layer, assigning 1 to all pixels that currently have Soil Class 1, 3, or 5, and 0 to all other pixels.

GIS software with well-developed raster capabilities, such as IDRISI software, allows the user to access a wide range of operations of this type that create a new layer through a local operation on an existing layer.

11.5.2 Focal Operations

Focal operations produce results by analysing each pixel in relation to its immediate neighbours. Some of these operations work on all eight immediate neighbours, but others focus only on the four neighbours that share a common edge and ignore the four diagonal neighbours. Sometimes we distinguish these options by referring to moves in the game of chess - the eight-neighbour case is called the queen's case, and the four-neighbour case is called the rook's case (Fig. 11.10).



Fig. 11.10: Two definitions of a raster cell's neighbourhood. In the rooks case only the four cells that share an edge are neighbours, but in the queen's case the neighbourhood includes the four diagonal neighbours

Among the most useful focal operations are the so-called *convolutions*, in which the output is similar to an averaging over the immediate neighbourhood. For example, we might produce a new layer in which each pixel's value is the average over the values of the pixel and its 8 queen's case neighbours. Fig. 11.10 shows a simple instance of this. The result of this operation produce an output layer that is smoother than the input, by reducing gradients, lowering peaks and filling valleys. In effect, the layer has been filtered to remove some of the variation between pixels and to expose more general trends. Repeated application of the convolution will eventually smooth the data completely.

Convolutions are very useful in remote sensing, where they are used to remove noise from images. Averaging can only be used if the values in each pixel are numeric and measured on continuous scales, so this operation would make no sense if the values represent classes of land, such as soil classes. But in this case it is possible to filter by selecting the commonest class in the neighbourhood, rather than by averaging (Fig. 11.11).

10	9	7	4	5
9	6	5	3	4
6	3	2	2	3

↓

8.5	7.5	5.7	4.7	4.0
7.2	6.3	4.6	3.9	3.5
6.0	5.2	3.5	3.2	3.0

Fig. 11.11: Application of a convolution filter. Each pixel's value in the new layer is the average of its queen's case neighbourhood in the input layer. The result is smoother, picking up more of the general trend. At the edge, where part of the neighbourhood is missing, averages are over fewer than nine pixels

Digital Elevation data, also referred to as digital elevation models (DEMs), or digital terrain data are digital representations of the shape of the Earth's surface. Typically digital elevation data consist of arrays of values that represent topographic elevations measured at equal intervals on the Earth's surface. Such data can be considered analogous to digital remote sensing images except that each pixel represents an elevation measurement rather than a brightness value. Digital elevation data can be manipulated, classified and analysed in many of the same ways that digital remote sensing data can be displayed as pictorial images.

Slope and aspect can also be used to estimate the pattern of surface water flow over a DEM (digital elevation model), a very useful operation in determining watershed boundaries and other aspects of surface hydrology. Each pixel's elevation is compared to those of its eight neighbours. If at least one of the neighbours is lower, then the GIS infers that water will flow to the lowest neighbour. If no neighbour is lower, the pixel is inferred to be a pit, in which a shallow lake will form. If this rule is applied to every pixel in a DEM, the result is a tree-like network of flow directions (Fig. 11.12), with associated watersheds.




Fig. 11.12: A digital elevation model, and the result of inferring drainage directions using the simple rule "If at least one neighbour is lower, flow goes to the lowest neighbour"

If the input layer is a digital elevation model, with pixel values equal to terrain elevation, a form of local operation can be used to calculate slope and aspect. This is normally done by comparing each pixel's elevation with those of its eight neighbours, and applying simple formulae (Burrough and McDonnell, 1998). The result is not the actual slope and aspect at each pixel's central point, but an average over the neighbourhood. Because of this, slope and aspect estimates always depend on the pixel size (often called the distance between adjacent postings) and will change if the pixel size changes. For this reason it is always best to quote the pixel size when dealing with slope or aspect in a GIS.

Another powerful form of local operation forms the raster equivalent of the buffer operation and also supports a range of other operations concerned with finding routes across surfaces. To determine a buffer on a raster, it is necessary only to determine which pixels lie within the buffer distance of the object. Fig. 11.13 shows how this works in the cases of a point and a line. But suppose we make the problem a little more complex, by asking for pixels that are within a certain travel time of a given point, and allowing travel speed to be determined by a new layer. Fig. 11.14 shows such a layer and the result of determining how far it is possible to travel in given numbers of minutes. The operation is also known as *spreading*. Of course this assumes that travel is possible in all directions, but this would be appropriate in the case of an aircraft looking for the best route across the Atlantic or a ship in the open ocean.




Fig. 11.13: The buffer operation in its raster form, for a point and a line (indicated by the white cells)




Fig. 11.14: Illustration of a spreading operation over a variable friction layer. From the central white cell, all cells in the black area can be reached in 10 minutes, and all cells in the grey area can be reached in 20 minutes

This method is used very frequently to find the best routes for power transmission lines, highways, pipelines, and military vehicles such as tanks — in all of these cases; the assumption of travel in all directions is reasonably accurate. Origin and destination points for the route are defined. A raster GIS is used to create a layer of travel speed, or friction. Because this is not uniform, the best route between the origin and destination is not necessarily a straight line. Instead, the best route is modelled as a series of moves in the raster from each cell to the most appropriate of its eight neighbours, until a complete route of least total time or least total cost is found (Fig. 11.15). There have been many applications of this GIS method over the past three decades by highway departments, power utilities, and pipeline companies.




Fig. 11.15: Finding the least-cost path across a variable-cost layer from a given origin to a given destination, using queen's case moves between raster cells. This form of analysis is often applied to find routes for power lines, pipelines, or new highways

The final operation considered here is the so-called *viewshed operation*. This also works on a DEM and is used to compute the area that can be seen from a specified point, by a person at that point. When more than one point is specified, the GIS could be used to determine the area visible from at least one point, or in some military applications it might be useful to determine the area invisible from all points. The viewshed operation is often used in planning. For example, if a new clearcut is proposed by a forest management agency it might be useful to determine whether the scar will be visible from certain key locations, such as public roads.

11.5.3 Zonal Operations

Raster data is represented in pixels, so the line objects and the area objects are modeled as collections of pixel values with no clear linkage between them. If we simply count the pixels assigned to Class 1 on a soil layer, for example, we get the total area of Class 1, not the separate areas of the individual patches of Class 1. Zonal operations attempt to address this by focusing on operations that examine a layer as a collection of zones, each zone being defined by contiguous pixels of the same class. Zonal operations can be used to measure the areas of patches or their perimeter lengths and the results are returned as new layers in which each pixel is given the appropriate measure.

11.5.4 Global Operations

Finally, global operations apply a simple manipulation to an entire layer, usually returning a single value. Examples include operations to determine the average of all pixel values or the number of pixels having a specific value, or the numbers of pixels having each unique value, or the total area having each unique value. For example, a global measurement of area on a soil layer with six classes might return six areas, plus the total area. Global operations produce summary statistics, so they are often the last operation in a sequence.