GEOGRAPHIC INFORMATION SYSTEMS AND DISAGGREGATE TRANSPORTATION MODELING

Michael F. Goodchild
National Center for Geographic Information and Analysis, and
Department of Geography
University of California
Santa Barbara, CA 93106-4060

ABSTRACT

Geographic information systems (GIS) have developed in response to a range of needs, and provide highly structured environments for working with particular classes of information. The environment and structure of a GIS is defined by its data model; the paper reviews the basic data models of GIS. The needs of discrete transportation modeling have led to a number of extensions of the basic data models. Further necessary extensions are identified. The paper ends with a discussion of appropriate strategies and future directions.

1 Introduction

Both geographic information systems (GIS) and discrete transportation modeling are burgeoning fields, characterized in recent years by rapid technological and scientific developments. GIS in particular is a large-scale coalescence of varied interests, ranging from the needs of utility companies for software and databases to manage geographically distributed and heterogeneous facilities, to the needs of geoscientists for analytic tools. The breadth of the field is amply illustrated by several recent texts (see, for example, Chrisman, 1997; Clarke, 1997; Maguire, Goodchild, and Rhind, 1991; Worboys, 1995). Definitions tend to be vague, and the use of the term clearly overlaps with many others, including digital cartography (Cromley, 1992), land information systems (Dale, 1991), and scientific visualization (Hearnshaw and Unwin, 1994).

Given the breadth of GIS, it is not immediately clear what it offers to discrete transportation modeling, or in what ways the two fields can interact to mutual benefit. The purpose of this paper is to explore this intersection in depth. It will thus be of interest to GIS specialists wanting to learn more about GIS support of discrete transportation modeling, and to transportation modelers interested in knowing more about the advantages and disadvantages of GIS. The paper also identifies important development directions for GIS if it is to play a larger role in discrete transportation modeling.

The next section of the paper explores the nature of GIS, focusing on its data models, and the major paradigms of GIS application. The third section provides a more general perspective on data models, and their importance in the construction and maintenance of software. Section Four identifies data model extensions that would be needed to make GIS more useful in discrete transportation modeling. The paper concludes with some comments on appropriate mechanisms for achieving such extensions.

2 GIS in theory and practice

A geographic information system is a digital computer application designed for the input, storage, manipulation, and output of geographic information; geographic information is defined as information referenced to specific locations on the surface of the Earth. Thus geographic databases represent the variation of phenomena over the Earth's surface. While there are potentially many ways to do this, in practice only a few are supported by the designers of GIS software, who tend to be driven by the needs of their users. Given the context of this paper, three broad classes of GIS representations are identified (Goodchild, 1992):

1. Schemes that represent continuous variation over the surface, one parameter at a time; these are termed *field* models. The best-known example of a field model in GIS is a representation of the elevation of the land surface; other similarly single-valued functions of geographic coordinates include surface temperature, or land ownership.

2. Schemes that represent collections of discrete point, line, or area features, and their associated properties; these are termed *discrete entity* models. This conceptualization is consistent with a view of the landscape as populated by discrete, potentially overlapping features such as lakes, houses, roads, contours, mountain peaks, or survey benchmarks.
3. Schemes that represent variation over a linear network embedded in the surface; these are termed *network* models. Such models are commonly used to represent transportation networks, and networks formed by surface hydrologic features.

The three schemes are not entirely mutually exclusive. A network (3) can be regarded sometimes as a collection of discrete linear entities (2). A single discrete entity (2) such as an area can be represented through the spatial variation in a binary parameter (1) which has the value of 1 inside the area and zero elsewhere, and collections of such entities can be treated similarly.

In practice, currently available GIS software offer six forms of field models. A field can be represented by storing its average value in each of a regular array of rectangular cells; its value at the center of each of a rectangular array of sample points (both of these options are included in the loose term *raster*); its value at each of a collection of irregularly spaced sample points; its average value in each of a set of non-overlapping, irregular areas; its value at the corner points of an irregular mesh of triangles; or by storing the locations of isolines (these last four are instances of structures loosely termed *vector*). Not all GIS offer all six, but all are represented by commonly available data sets.

While the field models are dominant in many GIS applications, notably in environmental science, where fields are used to represent the spatial variation of parameters such as topographic elevation, vegetation class, or mean annual temperature, they are clearly of less interest in transportation modeling. Attempts to model urban systems as fields of continuously varying parameters have been described by Angel and Hyman (1976) and many others, but are comparatively obscure. Numerous efforts have been made to model spatial interactions using continuous fields rather than discrete entities (e.g., Tobler, 1981), but they have not met widespread acceptance (e.g., Fotheringham and O'Kelly, 1989). Network models are clearly important, and many models of urban behavior are based on discrete entities (e.g., vehicles, houses, census tracts, shopping centers). Thus the emphasis in this paper will be on the latter two schemes: models of discrete entities, and networks.

The choice between alternative representations is driven by many factors. In science, a representation may be chosen because it is more accurate; and generally, greater accuracy is achieved by representations that capture greater detail, and thus occupy more storage space. A representation may also be preferred because it is closer to human conceptualization; for example, people may find it easier to understand a weather forecast expressed in terms of discrete entities (fronts, highs, lows) than continuous pressure fields, despite the fact that the models underlying such forecasts are predominantly field-based. This tension between scientific accuracy and intuitive understanding will be increasingly important in areas such as discrete transportation modeling, where human users must interact directly with scientific models and database structures in a range of environments. This point will be revisited several times in the discussion.

2.1 Discrete entity models

GIS databases store discrete entities as geometric structures according to their dimensionality. Points are stored as pairs of coordinates, lines as sequences of points assumed connected by straight lines (often known as *polylines*; connections by arcs of circles or spline functions are sometimes allowed), and areas as polygons. The term *vector* is commonly used to refer to all three, as well as to the four field models described earlier; vector data sets are thus characterized by the presence of coordinate pairs. Entities are grouped into classes, the members of each class being of uniform dimensionality (all points, all lines, or all areas) and being characterized by the same set of attribute variables. Each class may also be associated with a distinct set of files, tables in a database management system, and object methods. Support is normally provided for a range of attribute measurement scales, including names, numeric values, and dates. In essence, one can safely assume that a GIS can store everything that is known about any class of discrete entities.

In addition to entities and their attributes, a GIS may provide the framework for storage of relationships among entities. This is often achieved by storing the classes of entities and their attributes in relational tables, and using common keys between tables to store relationships such as connectedness, adjacency, or containment. For example, one might identify the county containing each case of a disease by creating two tables, one of counties as area entities, the other of cases as point entities, and storing the containing county's ID as an additional attribute of each case. This use of keys to store geographic relationships is sometimes called the *georelational* model. Note that it is effectively restricted to representing the existence of relationships, and does not allow the relationships to have their own attributes.

2.2 Network models

Clearly representations of connected linear networks are of great interest in transportation applications of GIS, since much information of interest in such applications is confined to linear networks, to an appropriate level of approximation. Thus network GIS models assume all activity is restricted to a linear network embedded in or lying on the Earth's surface. Representation occurs at two levels: representation of the locations of phenomena on the one-dimensional network, and representation of the locations of the one-dimensional network in the two-dimensional surface. For some purposes the latter can be ignored—as in many transportation models, outcomes may be the same however the network is embedded in the surface.

Other more restrictive assumptions are possible also. For many purposes one can assume that the linear network is homogeneous between nodes or junctions. Given no intermediate access, for example, traffic density must be constant between nodes. Thus many GIS representations allow only for attributes of entire links. These issues are discussed in greater detail in Section Four.

2.3 GIS application paradigms

The previous sections have identified the basic data models of GIS. With these, it is possible to create accurate digital representations of the contents of maps and Earth images, along with the kinds of information about geographic features that are normally found in statistical archives, rather than on maps. In this section, the applications that these basic GIS data models enable are broadly reviewed. Later, these application paradigms will be revisited in the context of discrete transportation modeling.

First, given the effectiveness of these data models in storing the contents of maps, it follows that GIS is a useful technology for *digital map production*. It contains the necessary tools and functions for input, editing, projection, scale-change, drafting, annotating, and symbolizing of paper maps. In essence, it enables everyone to be a map-maker (cartographer) whether or not they have formal training in the subject, or access to traditional tools. However this paradigm is clearly of less relevance to discrete transportation modeling.

Second, GIS is widely used for *inventory and management* of spatially distributed facilities. Utility companies use GIS to keep track of their installed assets, such as cables, pipes, and transformers, and their customer accounts. Transportation agencies use GIS to keep track of signage, pavement quality, rights-of-way, etc.

Third, GIS is a useful technology for *integration of data*. Its use of geographic location as a common reference allows otherwise unrelated data sets to be linked, analyzed jointly, combined, and used in supporting complex decisions. Because many data models and functions are supported, it is possible for integration to be achieved despite major differences in format, scale, projection, and data model. For example, a utility company might combine an image of an area with a vector representation of its underground pipes, in order to provide information to another agency planning to dig.

Fourth, GIS supports *spatial analysis*, or the manipulation of spatial data in order to extract information and insight. Technically, spatial analysis is defined (and distinguished from other forms of analysis) as a set of techniques producing results that are not invariant under continuous geometric distortion (Goodchild, 1987).

Finally, GIS supports *dynamic modeling*, using digital representations of data and processes to forecast impacts and evaluate scenarios, and to present results in the form of maps, charts, and tables. While elements of all of the five GIS application paradigms except the first are important in discrete transportation modeling, this last is clearly where the greatest potential overlap exists between the two fields.

3 The importance of data models

A data model is defined broadly as the set of entities and relationships used to create a representation of some real phenomenon or phenomena. Data modeling must occur whenever a representation is created, whether it be digital or not (the term *conceptual data model* is often used to distinguish this interpretation from other narrower ones). Thus a climatologist creates a data model when taking sample point measurements of the spatial variation in temperature over an area—in this case, what is conceptualized as a continuous field is being represented by a finite set of point measurements. If the data are later used to create a contour (isotherm) map, the data model may change from a set of points to a set of contour lines—or the result may be two representations, if the original point measurements are preserved as well. Later when the contours are digitized the representation may change again from a collection of smooth curves to the polylines used in a vector GIS database.

The range of data models allowed by the designer of a system defines the operations that are possible. For example, a system that can store a single rectangular raster layer can be used to display, or compute statistics summarizing the contents of the cells. By extending the data model to many such layers, it becomes possible to perform all of the standard functions of a raster GIS. A spreadsheet package, on the other hand, provides a set of functions around the data model of a simple table. Thus the versatility of a GIS is ultimately determined by the set of data models it enables—the most powerful GIS will be the one that implements the largest subset of the geographic data models identified earlier, assuming of course that the associated functionality is also provided. Similarly, the value of GIS to a field like discrete transportation modeling will depend on how effectively it can represent all of the data of relevance to the field—once the necessary data models are in place, it is comparatively easy for a system developer to add the appropriate functions.

The data models of GIS are highly specialized compared to those found in the larger world of database management systems (DBMS). The structures used by GIS to represent a continuous field using a mesh of irregular triangles, for example (the TIN model of GIS), are of little use to an airline reservation system. General-purpose DBMS use much less specialized data models, such as the relational model in which information is expressed in the form of linked tables. Relational DBMS are widely used in GIS, but only after the addition of specialized functions and structures to help the user implement the general model in the narrow spatial context. If they were not there, every user of GIS would have to solve the same basic problem, of how to express the entities and relationships of geographic data using the simple structures (tables and links) provided by the relational model. By providing standard solutions, vendors of those GIS that are based on relational DBMS add value to a general-purpose tool.

Having described both the special data models of GIS, and the general data models of standard DBMS, and the importance of data models in limiting the capabilities of software, the usefulness of GIS for discrete transportation modeling can now be examined.

4 Data models for discrete transportation modeling

In this section those data models most likely to be of value for discrete transportation modeling are examined, and extensions that have been made by GIS designers to deal with obvious deficiencies are identified. Cases that are clearly needed for discrete transportation modeling are discussed even if they are as yet not available in GIS. As noted earlier, the data models fall into the network and discrete entity classes with few exceptions.

4.1 Planar and non-planar networks

As noted earlier, a network can be conceived as a one-dimensional space embedded in a surface. The simplest model implemented in GIS represents each topological edge of the network as a polyline entity. Associated with each entity will be a set of attributes, conceived as the entries in one row of a rectangular table. This simple model is illustrated in Figure 1b, as a representation of the network shown graphically in Figure 1a. By adding a second table, it is possible to store the attributes of the network's topological nodes, as shown in Figure 1c. Figure 1d shows a scheme that includes storage of 'pointers' from edges to nodes, and from nodes to edges, to store the network's connectivity. Note that each edge has an inherent direction defined by the order of points in its polyline representation (and illustrated by the arrows in Figure 1a).

[Figure 1 here]

Several issues arise in deciding whether to store connectivity in this fashion. On the one hand, such connections should be computable, so the designer might choose to compute them 'on-the-fly' and save storage at the cost of additional computing. On the other hand it may not be possible to assume that the ends of edges coincide perfectly, especially if the data have been digitized, so there may be a risk associated with this approach. Moreover, one might want to distinguish between edges that meet in nodes, and are thus connected, and edges that cross in two-dimensional view without intersection (overpasses or underpasses). *Planar* data models force nodes at all intersections (Figure 2a), whereas *non-planar* do not (Figure 2b).

[Figure 2 here]

4.2 The turn-table

The distinction between planar and non-planar provides only a partial solution to the problem of connectivity. In modeling transportation, it may be necessary to include extensive information on the ability to connect from one edge to another. For example, drivers may face turn restrictions, or trucks may be limited by turning radius—both of these situations require more than the simple ability to represent the existence of a crossing at grade or an overpass.

To deal with this problem, the basic GIS data model has been extended with the addition of a new structure called the *turn-table*. For every ordered pair of edges incident at a node, a row of attributes in the table gives appropriate properties of the turn, together with links to the tables containing the attributes of the edges. In this way, a data model that imposes the planarity restriction noted above can still represent overpasses successfully, by preventing turns. Figure 3 shows a turn-table for the layout used in Figure 2.

[Figure 3 here]

4.3 Dynamic segmentation

Thus far, the data models presented have been capable of representing only those properties of networks of importance in discrete transportation modeling that can be conceptualized as attributes of entire edges or nodes. Some of the earliest attempts to create digital representations of networks were made by the DIME project of the U.S. Bureau of the Census in the 1960s (Coppock and Rhind, 1991), and later extended by TIGER. But the primary aim of both projects was the administration of the Census—applications of digital representations of networks for transportation were not important to the designers of these databases.

In this edge/node (or *link/node*) structure the only feasible way to represent a property or attribute of part of an edge is to add intermediate nodes. If the property refers to a point, it can then be represented as an attribute of a new node; if it extends for some fraction of the edge, then two new nodes must be introduced and the edge split into three, one of which receives the attribute. As links are split and nodes added the respective tables in the database must be extended and updated, and eventually the database will become unmanageable.

The response of many GIS designers has been to add a new structure, implementing *dynamic segmentation*. The basic edge/node structure is preserved, but above it is added a structure representing two types of discrete entities—zero-dimensional and one-dimensional—located at arbitrary locations on the network. These entities are given their own attribute tables, and will be termed here *network points* and *network segments* respectively. In effect, dynamic segmentation provides a mechanism for storing 0-D and 1-D discrete entities located in a one-dimensional space (the network). and for describing how that network is in turn embedded in the two-dimensional geographic space. Figure 4 shows a simple illustration of the principle. In some applications it may be desirable to allow a property to vary along a network segment, rather than remain constant. For example, change of elevation could be described by storing gradients over network segments, if the linear approximation was acceptable.

[Figure 4 here]

4.4 Route and milepost schemes

As noted earlier, the edge/node scheme was made popular in early GIS by the DIME and later the TIGER databases. In an urban area dominated by a gridiron street pattern each edge is of similar length, and most intersections are at grade. In rural areas, however, intersections can be very infrequent, with great variation in edge length. It may be more difficult to define 'intersection' without reference to the classes of roads, and construction of a new road may force a redefinition of the basic topologic structure. Even in urban areas, there may be problems of definition in the case of laneways, or the complex intersection geometries of freeways or traffic circles. All of these arguments point to a basic weakness and instability of the edge/node structure.

Moreover, however accurate it may be, the edge/node structure and its dependence on intersections may not represent the way we think of the network. We name roads and streets independently of intersections, causing redundancy and the potential for error when names have to be entered repeatedly for each edge of a long street. Highway departments often report locations by identifying a route and a distance from the route's beginning, not from the nearest intersection. As a general principle, computer systems are easiest to use when the structures they present to the user are as close as possible to those of the user's own conceptualization (e.g., Mark and Frank, 1991).

For all of these reasons *route/milepost* schemes have been proposed (e.g., Dueker and Vrana, 1992), and in some cases implemented, as alternatives to edge/node schemes. An example is shown in Figure 5. Its basic entities are routes, with associated attributes, and intersections. By linking the dynamic segmentation scheme to this structure rather than the edge/node structure, it is possible to store features that extend through nodes. Indeed, it seems very unlikely that one would want to use a route/milepost scheme without an associated dynamic segmentation structure.

[Figure 5 here]

4.5 Lanes

Thus far, discussion has been limited to networks as linear systems with no transverse structure. Although certain information about the transverse structure, such as the number of lanes, or the existence of a median, might be stored as attributes of edges or network lines, it is impossible to store detailed information about individual lanes, or connectivity at the lane level. For example, there is currently no way to disaggregate a turn-table to store turn restrictions that are specific to lanes.

To the modern driver, instructions about use of lanes are increasingly important. For example, it may be useful to know whether a freeway exit ramp occurs from the leftmost or the rightmost lane; whether all vehicles in a given lane are forced to exit; whether only the rightmost lane is allowed to turn; when to change lanes on approaching an intersection; etc. Instructions about lanes may be more useful than instructions about edges and nodes, in some driving environments. It follows that human behavior may be more easily modeled using lane-based representations. Moreover, in rail networks similar points can be made about the use of sidetracks, and the connectivity of the network at the track level. All of these arguments lead to a basic dissatisfaction with linear representations.

On the other hand, the level of geometric detail required for effective representation of lanes and tracks is clearly much higher. While 50m accuracy may be sufficient to locate a vehicle on a road, better than 5m will be needed to locate to the lane level. Such accuracies are well beyond the capabilities of many of the currently available network databases. For example, TIGER was developed by combining maps at 1:24,000 and 1:100,000, with positional accuracies of roughly 12m and 50m respectively. Achievement of better than 5m accuracy with GPS requires the use of differential techniques, and a high quality of geodetic control.

Fohl et al. (1996) describe an effective compromise. Each lane is represented as a distinct entity, with its own connectivity with other lanes, but its geometry is obtained from the standard linear geometry of the road. No attempt is made to store the relative positions of lanes, but the structure does identify such topological properties as adjacency, and the order of lanes across the road. Figure 6 illustrates the basic principles of the structure in a simple example.

[Figure 6 here]

With this structure, it is possible to compute and output driving instructions at the lane level.

4.6 Off-network travel

While most travel is effectively modeled by a network, certain modes present problems because of the possibility of travel off the network. Vehicles leave the linear road network when they enter parking lots, or travel routes that are not represented in the network database. Boats may follow an infinite number of tracks across open water; aircraft may in certain circumstances deviate from a linear network of flight paths. To model these situations, it is necessary to combine the linear network representation with a digital representation of continuous 2-D spaces, using one of the discrete entity or field models.

4.7 Flows

Transportation modeling may require the representation of flows, for example as counts of vehicles traveling between defined origins and destinations. From a data model perspective, several distinct cases can be identified:

1. Flows allocated to entities in the data model. Flows may be expressed as attributes of edges, and conservation laws may be applied at nodes. Alternatively, flows may be expressed as attributes of network lines. These representations will support analysis of congestion on edges or routes in the network, and may be the product of traffic assignment models or direct observation.
2. Flows between origins and destinations (square case). Flow data are available between pairs of places, in both directions. These places will normally be zones, represented as discrete area entities. However, the assignments of flows to various network elements are not known.
3. Flows between origins and destinations (rectangular case). As (2), but the sets of origin and destination places are different. For example, origins may be neighborhoods (discrete area entities) and destinations may be shopping centers (discrete point entities).

These three cases are illustrated in Figure 7.

[Figure 7 here]

While it is easy to represent (1) in the form of attributes of entities, the GIS data models described earlier provide no structures suitable for storing (2) or (3). Goodchild (1987) has called these examples of a general structure termed an *object pair*. The need for this structure occurs whenever attributes must be assigned to qualify or characterize a relationship between two geographic objects. In some cases, such as (1) above, these attributes can be assigned to other geographic objects, such as edges. In other cases, such as (2) and (3), no suitable geographic object exists, so an additional object pair structure must be supported. The turn-table illustrated in Figure 3 is also an example of an object pair structure. Although it is possible to store such data as a table in a relational database (each row of the table would contain an origin entity ID, a destination entity ID, and measures of flow or other property), no GIS currently recognizes the structure as an explicit form of geographic data model, or provides the necessary functions for processing, displaying, linking, or analyzing such data.

4.8 Complex paths

It is possible to store a path in an edge/node structure by storing an ordered sequence of nodes, or an ordered sequence of edges. Similar methods could be used to store a path in a route/milepost structure. As in the previous section, however, no GIS currently provides facilities for storing and processing such entities in its basic data model. Moreover, there may be situations in which such paths may have attributes that are true only of the whole, and not of each part of the path, and thus cannot be stored as attributes of the path's component edges. For example, it may be necessary to store a number of bus routes, together with associated attributes like route number, schedule, and ridership. A delivery company might need to store the delivery routes taken by its vehicles each day. One might also want to store the fact that the traffic lights along a path were synchronized—an attribute that is meaningless if stored separately for each edge of the path, but may be very important for accurate modeling of behavior. The inability to store such paths and their attributes, whether as inputs or as the results of modeling, is clearly an impediment to the use of GIS for transportation modeling.

4.9 Temporal change

Early GIS data models were intended for the representation of maps and images, and so paid no attention to temporal change. Recent research has provided many solutions to the representation of temporal change (e.g., Langran, 1992), but few have been implemented in current GIS. Thus it is straightforward to store the congestion at different times of day as a series of attributes in a table of edges; but the GIS will provide no tools for processing the series as an inherently ordered temporal sequence. Similarly, the user is free to devise ways of representing time-dependent turn restrictions in turn-tables, but the correct interpretation of the resulting attributes is not aided in any way by standard system features.

Discrete transportation modeling is likely to require the ability to store locations of discrete elements, such as vehicles, within the GIS data model. In their work on diurnal travel behavior in Halifax, Goodchild, Klinkenberg, and Janelle (1993) used two distinct ways of representing the space-time behavior of individuals and integrating them with GIS data models (see Figure 8):

1. Activity events (Figure 8a). These are periods during which an individual is engaged in a single activity. Associated attributes include the person’s ID, start time, stop time, activity code, location at start, location at stop. Each event is thus associated with two point locations, which will be identical if the activity is static. In the case of travel activities it may be possible to assign the individual to a transport network and infer intermediate locations, but the assumptions that would have to be made are in most cases unreasonable.
2. Space-time matrices (Figure 8b). By assigning activity events to time periods and discrete area entities, it was possible to assemble a three-dimensional matrix of areas by time periods by activities, giving the numbers of individuals in each cell. The information in these matrices could be treated as attributes of the area entities and mapped; or used as input to three-dimensional analysis to identify aggregate patterns of space-time behavior. In many current GIS it is possible to store them in relational tables and link them to area entities, but as in the previous section no GIS provides the structures needed to recognize time intervals as an ordered sequence.

Conclusion

Discrete transportation modeling makes use of a wide variety of data types, ranging from aggregate flows from origins to destinations, to the behavior of atomic elements in space and time. GIS data models originated in the need to represent the contents of static maps, but have gradually been extended to handle the special needs of spatial analysis and modeling. These include many of the kinds of analysis and modeling needed in transportation.

In this paper, we first identified the core GIS data models, and identified two of the three general classes as important to discrete transportation modeling. The significance of a data model lies in its role in defining the ability of a system developer to add functions, by providing the structures needed to store the essential data. When these structures are highly specialized to a given application, the user is provided with a well-defined environment adapted to his or her particular needs. When the structures are general, the onus is on the user to adapt the framework to the needs of the application. Seen from this perspective, the data models supported by GIS are extensive, but in general much more highly specialized than those underlying spreadsheets, the statistical packages, or relational database management systems. While virtually any form of information can be represented in an RDBMS (though there is an extensive literature on the inadequacies of the relational model), the specialized structures of a GIS allow the system's designers to provide a wide range of sophisticated techniques and tools for geographic data, provided the data obeys the rules and constraints inherent in the data model. At the same time their specialization makes it difficult to add extensions, such as those identified in this paper as appropriate for discrete transportation modeling.

Section Four showed that discrete transportation modeling has requirements for data types that go well beyond those commonly available in current GIS. Many GIS now support dynamic segmentation, route/milepost network representation, and turn-tables. Recent research has identified extensions for storage of information at the lane level. But the discrete transportation modeler is still faced with significant gaps, notably in the representation of flow matrices and paths. To deal with them, three options are currently available:

1. Adopt a GIS that includes an RDBMS, and build the tables necessary to support the alien data types. The GIS will provide no special tools for management or display of these data types, and it may be necessary to export data to other environments for specialized analysis.
2. Adopt a more general software environment, such as an RDBMS or a statistical package, which may provide a broader range of general-purpose analysis tools, but will offer no specialized tools for geographic data.
3. Partition the database, storing appropriate elements in a GIS and others in a more general environment. This option may be impractical, because of problems of interoperation between the two environments.

In addition, a fourth alternative may be to adopt a GIS that is suited to data model and software extensions. Although these are long-term objectives of the Open GIS Consortium (http://www.opengis.org), it may be some time before they are fully realized.

In practice, the choice between (1) and (2) is complex and difficult, and may be dictated by extraneous factors such as familiarity with GIS, or disciplinary background. (1) will be favored in cases where the problems of management of geographic data, including preprocessing and retrieval from large stores, are substantial; where the display capabilities of the GIS are useful; and where the export of data to specialized analysis or modeling environments is comparatively straightforward.

Although GIS has been presented here as a comprehensive computing environment for handling geographic information, several significant niches have emerged within its broad umbrella. Certain GISs already provide better support for discrete transportation modeling than others, by providing more of the extensions identified above. One might reasonably ask, therefore, whether in time a distinct class of software environment will emerge—perhaps a DTMIS (discrete transportation modeling information system)—with data models designed to be of greatest value to this application. This will depend of course on the potential size of the niche, and its homogeneity. But the historical factors that have led to the emergence of niche environments in the past, GIS among them, have been weakened in recent years by a number of technical developments (Sondheim, Gardels, and Buehler, 1997). Increasingly, today's software environments make it possible to process one previously alien data type entirely within the environment of another. In the future, we are likely to see much weaker cleavages between software environments, and the cleavages that exist will be much less driven by data model distinctions. The result may be a trend towards a much lower granularity of software, defined by function rather than data model, and supported by much easier integration and linkage between modules. But it will be some time before these developments impact the software environments available to the discrete transportation modeler.

References

Angel, S., and G.M. Hyman (1976) *Urban Fields: A Geometry of Movement for Regional Science*. London: Pion.

Chrisman, N.R. (1997) *Exploring Geographic Information Systems*. New York: John Wiley and Sons.

Clarke, K.C. (1997) *Getting Started with Geographic Information Systems.* Upper Saddle River, NJ: Prentice Hall.

Coppock, J.T., and D.W. Rhind (1991) The history of GIS. In Maguire, D.W., M.F. Goodchild, and D.W. Rhind (editors) *Geographical Information Systems: Principles and Applications.* London: Longman Scientific and Technical, Vol. 1 pp. 21–43.

Cromley, R.G. (1992) *Digital Cartography*. Englewood Cliffs, NJ: Prentice Hall.

Dale, P.F. (1991) Land information systems. In Maguire, D.W., M.F. Goodchild, and D.W. Rhind (editors) *Geographical Information Systems: Principles and Applications*. London: Longman Scientific and Technical, Vol. 2 pp. 85–99.

Dueker, K.J., and R. Vrana (1992) Dynamic segmentation revisited: a milepoint linear referencing system. *Journal of the Urban and Regional Information Systems Association* 4(2): 94–105.

Fohl, P., K.M. Curtin, M.F. Goodchild, and R.L. Church (1996) A non-planar, lane-based navigable data model for ITS. In M.J. Kraak and M. Molenaar (editors) *Proceedings, Seventh International Symposium on Spatial Data Handling, Delft*, August 12-16, pp. 7B.17-7B.29.

Fotheringham, A.S., and M.E. O'Kelly (1989) *Spatial Interaction Models: Formulations and Applications*. Dordrecht: Kluwer.

Goodchild, M.F. (1987) A spatial analytical perspective on geographical information systems. *International Journal of Geographical Information Systems* 1(4): 327–334.

Goodchild, M.F. (1992) Geographical data modeling. *Computers and Geosciences* 18(4): 401-408.

Goodchild, M.F., B. Klinkenberg, and D.G. Janelle (1993) A factorial model of aggregate spatio-temporal behavior: application to the diurnal cycle. *Geographical Analysis* 25(4): 277-294.

Hearnshaw, H.M., and D.J. Unwin, editors (1994*) Visualization in Geographical Information Systems*. Chichester: John Wiley and Sons.

Langran, G. (1992) *Time in Geographic Information Systems*. London: Taylor and Francis.

Maguire, D.W., M.F. Goodchild, and D.W. Rhind, editors (1991*) Geographical Information Systems: Principles and Applications*. London: Longman Scientific and Technical.

Mark, D.M., and A.U. Frank, editors (1991*) Cognitive and Linguistic Aspects of Geographic Space*. Dordrecht: Kluwer.

Sondheim, M., K. Gardels, and K. Buehler (1997) GIS interoperability. In P.A. Longley, M.F. Goodchild, D.J. Maguire, and D.W. Rhind, editors, *Geographical Information Systems: Principles, Techniques, Management, Applications*. Cambridge: GeoInformation International, in press.

Tobler, W.R. (1981) A model of geographical movement. *Geographical Analysis* 13(1): 1–20.

Worboys, M.F. (1995) *GIS, A Computing Perspective*. London: Taylor and Francis.

Figure 1: Representation of the edges and nodes of a network. (a) graphical representation, (b) a simple edge table, (c) a simple node table, (d) pointers added to edge table and node table to represent connectivity.

(a)

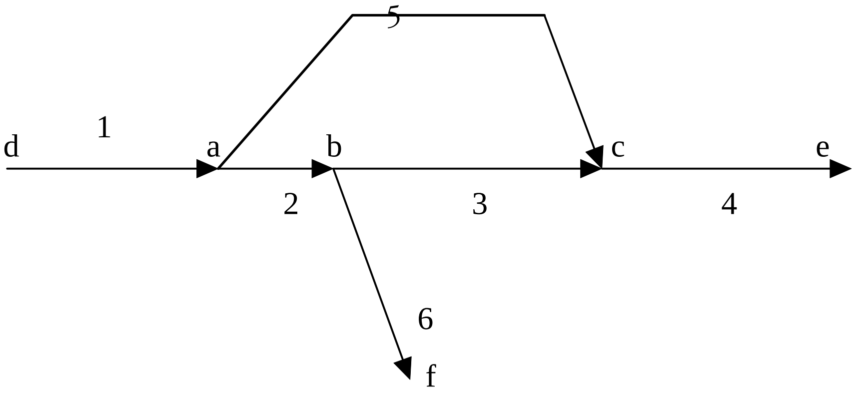

(b)

Edge ID	Street name	Lanes
1	High Street	2
2	High Street	4
3	High Street	4
4	High Street	2
5	River Way	2
6	Hill Street	2

(c)

Node ID	Stop light?
a	y
b	y
c	n
d	n
e	n
f	n

(d)

Edge ID	Street	Lanes	From node	To node
1	High Street	2	d	a
2	High Street	4	a	b
3	High Street	4	b	c
4	High Street	2	c	e
5	River Way	2	a	c
6	Hill Street	2	b	f

Node ID	Stop	Edge list
a	y	1,2,5
b	y	2,3,6
c	n	3,4,5
d	n	1
e	n	4
f	n	6

Figure 2: Illustration of (a) planar and (b) non-planar network approaches to representation of an overpass.

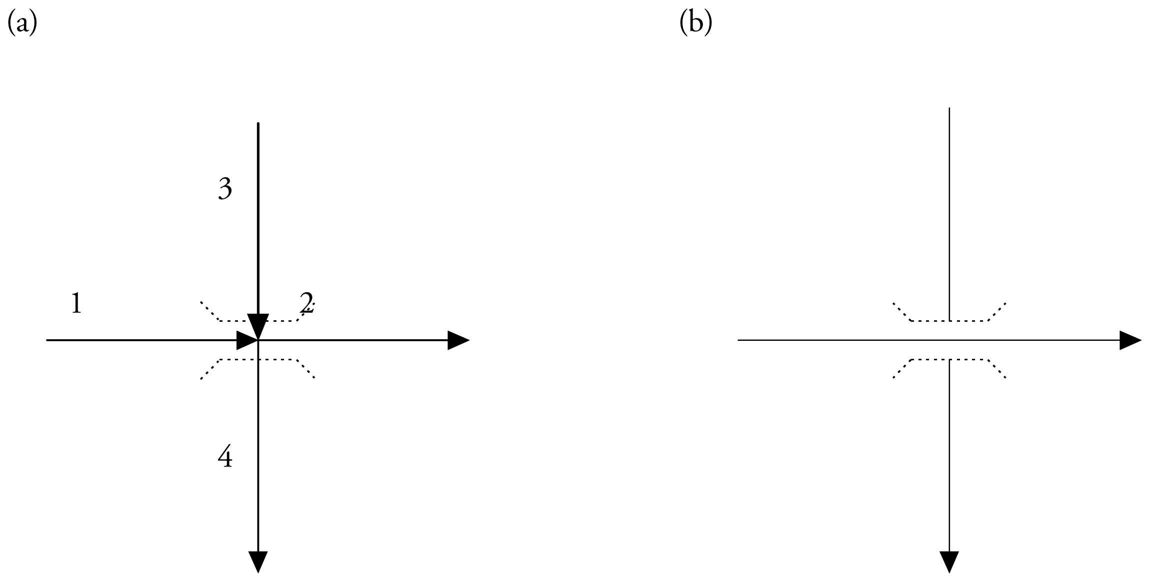

Figure 3: Use of a turn-table to prohibit turns, allowing non-planarities (e.g., Figure 2) to be captured in planar representations.

From edge	To edge	Turn?
1	3	n
1	2	y
1	4	n
2	1	y
2	3	n
2	4	n
3	1	n
3	2	n
3	4	y
4	1	n
4	2	n
4	3	y

Figure 4: The concept of dynamic segmentation: example structures of (a) network points and (b) network segments.

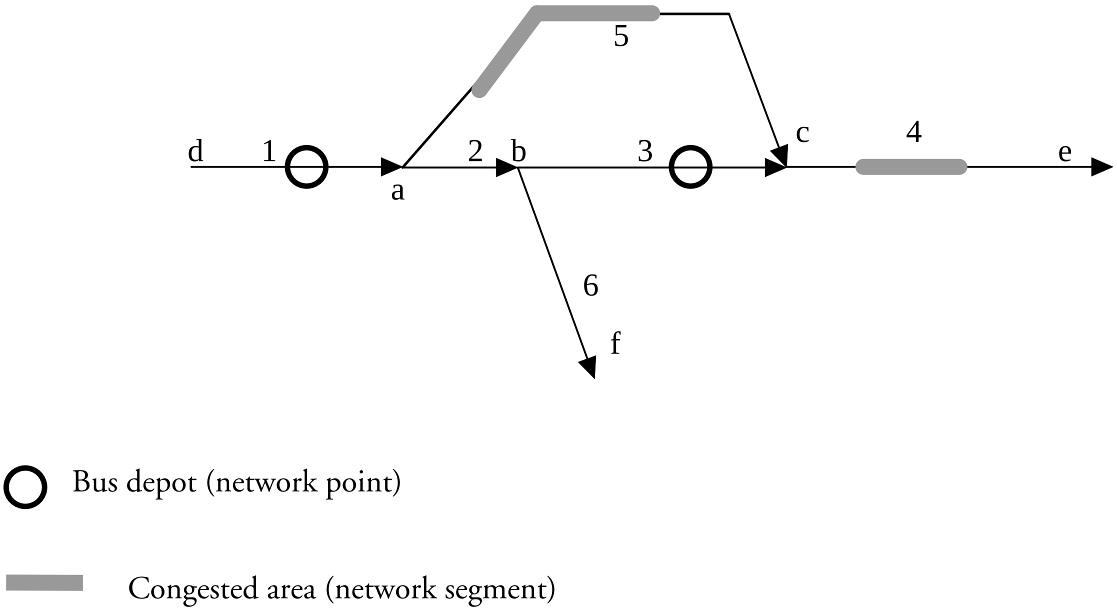

(a)

Edge	Distance from edge start	Feature
1	0.4	Bus depot A
3	0.8	Bus depot B

(b)

Edge	Distance from edge start to start of feature	Distance from edge start to end of feature	Level of congestion
4	0.3	0.8	High
5	0.5	1.9	High

Figure 5: Example of a route/milepost scheme, with associated structures: (a) graphic representation, (b) route table, (c) intersection table, (d) network point table.

(a)

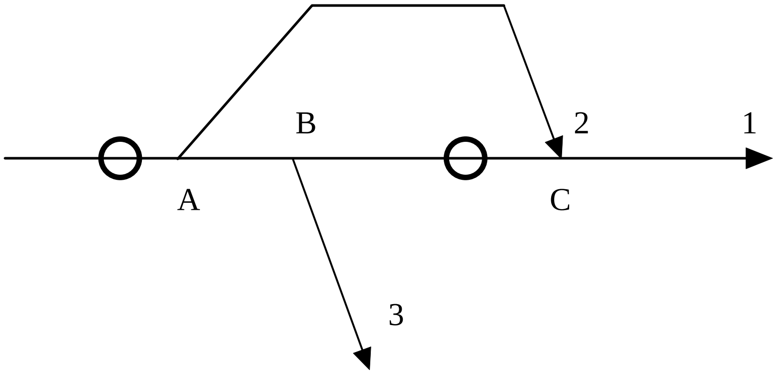

(b)

Route ID	Street name
1	High Street
2	River Way
3	Hill Street

(c)

Intersection	First route	Milepost	Second route	Milepost	Stop light?
A	1	0.8	2	0.0	y
B	1	1.6	3	0.0	y
C	1	3.0	2	2.4	n

(d)

Route	Milepost	Feature
1	0.4	Bus depot A
1	2.4	Bus depot B

Figure 6: Example of a lane-level structure. (a) graphic depiction of the street layout; (b) table providing information about each lane, including the lane's position relative to edges and nodes (see Figure 4 and the concept of dynamic segmentation of edges); (c) table of allowed turns between lanes (additional attributes of turns, such as time restrictions, could be added to this table).

(a)

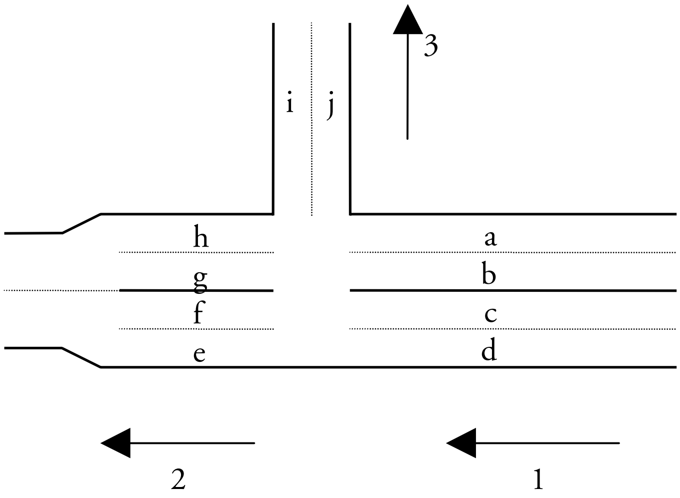

(b)

Lane	Edge	Start	End	Order from edge	Travel direction
a	1	0.0	1.2	2	forward
b	1	0.0	1.2	1	forward
c	1	0.0	1.2	1	back
d	1	0.0	1.2	2	back
e	2	0.0	0.4	2	back
f	2	0.0	1.2	1	back
g	2	0.0	1.2	1	forward
h	2	0.0	0.4	2	forward
i	3	0.0	1.4	1	back
j	3	0.0	1.4	1	forward

(c)

From edge	To edge
a	h
a	j
b	g
e	d
f	j
h	g
i	h
f	c

Figure 7: Three distinct flow data models. (a) flows assigned to edges; (b) OD flows, the same set of objects as both origins and destinations (square case; journeys to work between census zones); (c) OD flows, origin and destination objects distinct (shopping flows from neighborhood zones to shopping centers).

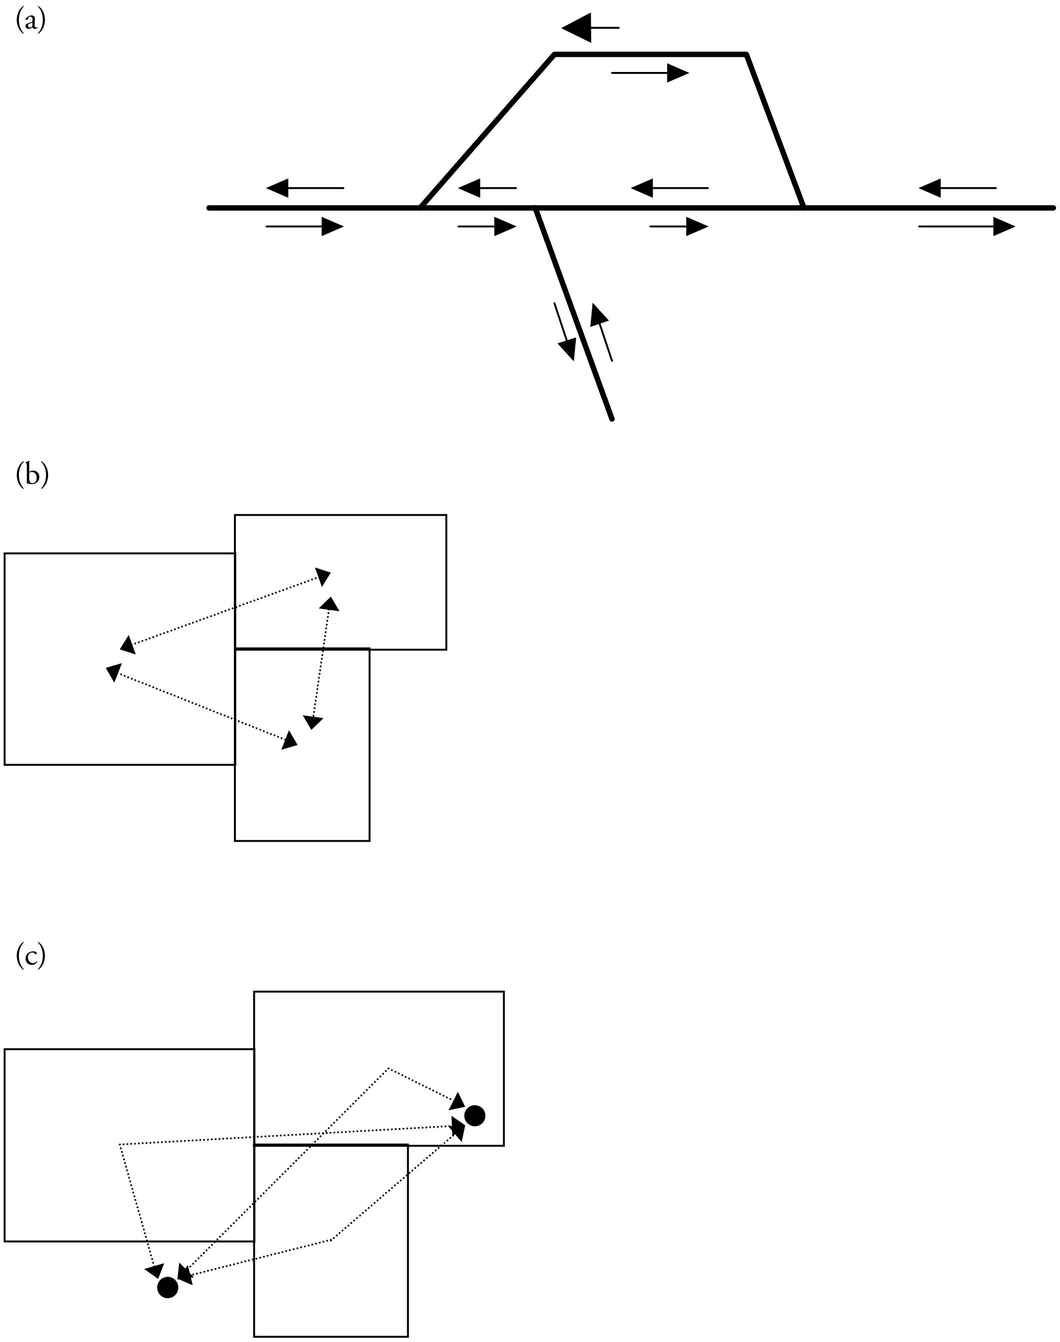

Figure 8: Two methods of representing the space-time behavior of individuals. (a) activity events; (b) a space-time matrix.

(a)

Person	Start time	Stop time	Activity	Location at start	Location at stop
1	0000	0700	sleep	7854,6743	7854,6743
1	0700	0800	eat	7854,6743	7854,6743
1	0800	0845	commute	7854,6743	6302,7221
….					

(b)