24

GIS interoperability

M SONDHEIM, K GARDELS, AND K BUEHLER

The term 'interoperability' has a range of meanings, but all focus on the ability to move easily from one system to another. This chapter reviews various strategies that may be used to achieve degrees of interoperability, including common exchange formats, conversions, common interfaces, and common database models. Geographical data modelling lies at the core of the interoperability issue, since agreement is needed on the representational framework. The chapter also covers issues of geographical metadata, common catalogues, and other tools to support search and retrieval across systems.

1 INTRODUCTION

Interoperability has been a goal of the computer industry for years and is essentially one view of the push for open systems. The term 'open systems' usually implies the intention to adhere to vendor-neutral computing standards, with the added benefit of producing a more level playing field among software and hardware companies. As a means of achieving open systems, interoperability has centred on common communications infrastructures, application programming interfaces in the public domain, and a common architecture for defining objects and transporting them across networks. Open computing and interoperability are as vital to geographical information systems as to any other area of information technology.

In the GIS arena interoperability has been a serious concern since the late 1970s. Problems with incompatible computing environments have been compounded by the inherent complexity of geographical information and the many ways in which it can be modelled. Nevertheless, some interoperability schemes are viable today, and more powerful ones are under development. This chapter reviews the general notion of interoperability, strategies for GIS interoperability, geospatial data models, and high-level concepts about information sharing.

2 GENERAL NOTION OF INTEROPERABILITY

Where applications require more than the operations available on a single desktop and dataset, data users and providers must achieve a new set of objectives regarding heterogeneous data and distributed computing resources.

- Data producers must ensure that their data are readily accessible and understandable to potential data consumers.
- Users must be able to identify and locate relevant information, and know whether a given set of data is germane to their work.
- Queries to dispersed sites must be formulated and processed in a manner meaningful to both data server and application client.
- Geodata from one source must be capable of being integrated with data from another, in terms of both structure and semantics.
- Display and analytical functions must be associated with particular data models and made available to the requester.

In some situations, meeting these requirements is relatively straightforward: a distributed application may be designed in a top-down fashion across an enterprise, based on common technology and a common semantic framework. The enterprise standardises on a given technical infrastructure, typically provided by a single vendor or systems integrator. Individuals in the enterprise use a distributed database and associated applications, built on a limited set of data models and targeting particular business needs. This approach of common semantics and technology may be scaled up to include a group or community of organisations as long as all of the participants are willing and able to adopt the new system. Although such scenarios meet the primary goals of interoperability, they are usually described as integrated systems, rather than interoperability solutions.

Interoperability usually refers to bottom-up efforts (Litwin et al 1990), neither imposed by a central authority nor driven by a single application. The systems and data models found among the users are heterogeneous, having been developed independently of one another; consequently, systems re-engineering may be required to meet basic interoperability requirements. Two significant challenges must be met:

1. the autonomous systems must be able to exchange data and to handle queries and other processing requests;
2. they must be able to make use of a common understanding of the data and requests.

The first requirement implies that a common set of services must be universally available and accessible through network communications. The second dictates that a common formal language and a common model representation be used (UCGIS 1996). The semantics of the data, as conveyed through the language and model, are either used directly by each system, or are transformed to local semantic constructs which can be interpreted directly by the users' applications.

3 GIS INTEROPERABILITY STRATEGIES

3.1 Direct translation

Over the last two decades, much practical interoperability work has boiled down to exchanging geographical and computer-aided-design data by using bi-directional translators operating on vendor-specific data formats. The architecture of most translators amounts to a data reader, a correlation table defining the correspondence between input and output data types and a writer. The correlation table, which may be hard-wired into the translator, defines how given data types and values in the input stream should relate to data types and values in the output stream (Figure 1). Consequently, where the input and output data models are dissimilar, a good match may not be possible, resulting in loss of information. Often such translators are most successful when tailored for specific datasets.

A more powerful approach to translation is shown in Figure 2. In this case the software maintains an internal data model which is semantically more rich than either the input or output models. The input data types and values are mapped to types and values of this internal model. The transfer data types play essentially the same role as a common format; however, in this case, the intermediary form is a transient in-memory representation. Once the data are in the transfer data types, they may be redefined if desired through a series of transformations and geoprocessing steps available in the translation software. The software may even support feature redefinition in which a number of input features define a single output feature, or vice versa. Practically, it becomes possible to consider very different input and output models and to infer matches impossible with direct correlation. This kind of processing may be termed smart translation or semantic translation.

Consider forest stand data in a given system with the attributes stored partly in labels and partly in column-aligned ASCII tables, and with the area for each stand represented by a point and a set of bounding arcs which are also shared by the adjacent areas. Now assume that the data are desired in another system where each area is represented as an independent polygon with no inside point, and with all attributes in a dbf file (a common standard file format for databases). Additionally, users may want to simplify the line work and amalgamate small polygons with their most similar neighbours. It may also be desired to take a series of forest stand files, turn them into a seamless coverage, and then clip the result to a given watershed boundary. To carry out such translations is not possible with the first architecture, but is with the second.

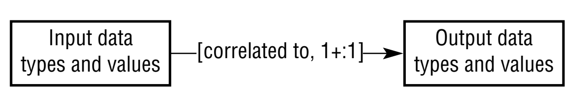

Fig 1. Simple translation through correlation.

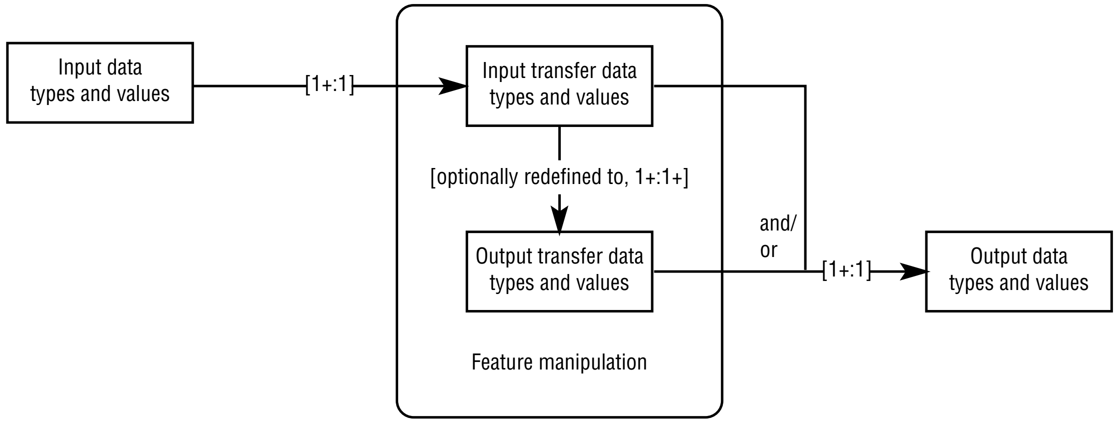

Fig 2. Semantic translation through feature manipulation and correlation.

Regardless of architecture, direct translation depends on the ability to read and write commercial formats. Some vendors have made their format specifications publicly available, but others have not. In the latter case, reverse engineering may be possible, but long-term maintenance of the translator may be costly because of undisclosed format changes by the vendors as their products evolve. Another option for vendors is to provide software interfaces (application programming interfaces, APIs) to their data management applications, as discussed in section 3.3.

3.2 Common transfer format

A variation of simple translation is the use of a basic format such as DXF or DLG (for a more detailed discussion of exchange standards see Salgé, Chapter 50) as part of a two-step process. Data in System A format are first translated to the intermediate format and from there a second translation converts them to the System B format. As used, these basic formats represent a lowest common denominator between the two end systems, possibly resulting in information loss. Additional limitations are that there may be no means of including the definitions of the types of data within the file, of handling metadata, of addressing updates, or of translating some kinds of data at all.

A number of efforts have centred on developing loss-less file exchange standards, based on non-proprietary file formats. Some of these standards have been aimed at resolving specific problems, such as a government agency receiving topographic data from private contractors in a neutral format (IEF in Israel or MOEP in British Columbia). Some have been designed to work within given disciplines or areas of endeavour. The International Association of Geodesy and other groups working with the Global Positioning System (GPS: see Lange and Gilbert, Chapter 33) recommend use of the receiver independent exchange format (RINEX) for the exchange of GPS data for post-processing. The International Hydrographic Organisation has developed S–57 as a transfer standard for digital hydrographic data, with the long-term intention of replacing traditional nautical charts around the world with electronic navigation charts supplied in S–57 format.

Other loss-less file-based efforts have been much more general in intent and have included sophisticated data models for geographical data. A clean separation is made between the logical model and the data encoding, the physical representation in a file. The Spatial Data Transfer Standard (SDTS) is a standard from the US government (US Geological Survey 1996) that provides a logical data model and a physical encoding (ISO 8211). Sponsored by the Digital Geographic Information Working Group, associated with NATO, the Digital Geographic Information Exchange Standard (DIGEST) is a collection of file exchange standards for different types of data (Digital Geographic Information Working Group 1995a). It has many similarities to SDTS and also refers to ISO 8211 for encoding. The Spatial Archive and Interchange Format (SAIF) is a Canadian standard based on an extensible, object-oriented paradigm and a formal modelling language (Geographic Data BC 1995). SAIF makes use of a zipped ASCII encoding which also allows for the direct inclusion of binary data. SDTS, DIGEST, and SAIF all provide support to developers through APIs to data files in their respective data encodings (see section 3.3).

Using the approach of a common format, each type of data holding is mapped to a common model, as manifest in a common format. Figure 3 shows some typical situations where the geometry and traditional attributes may be stored together or separately. As well, different data models may exist in each of these cases; for example, Model 2 and Model B may be quite dissimilar, even though they have similar implementations. The figure also shows the common model as implemented in a single file of a given format. Although this may be the case, an alternative is a series of files in given formats associated with particular types of data.

Use of a common format typically involves two operations: translating a file in a given format to an intermediary file in a common format, and then translating that intermediary file to a third file in the desired format. Only one step is required, of course, if the data are available in the common format, or if the common format is the desired format. Carrying this argument one step further, if the start and end formats are identical then translation is never required. In fact, there are many situations where Format X to Common Format to Format X is worthwhile, if the translation software provides quality assurance benefits or simply redefinition within the same format.

Work has also been underway to extend the data-centric view of files to a more processing-centric view. A technical committee of the Comité Européen de Normalisation (CEN) is examining geographical information (CEN/TC 287 1996) with the intent of developing a transfer method based on a data model and encoding, as with the standards described above. However, it also makes use of EXPRESS (ISO 10303-11), a relationally-oriented language appropriate for handling traditional attributes. Through EXPRESS, query and update operations may be supported. Another committee (ISO/TC 211) under the auspices of the International Organisation for Standardisation (ISO) is working on an approach similar to that of CEN (ISO/TC 211 Secretariat 1996), but with the intention to include operators and services to enable logical model to logical model transformations. As a separate endeavour, the FMEBC (Geographic Data BC 1996) is a sophisticated translation software package associated with the SAIF standard and specifically SAIFLite (a pared-down version of SAIF for operational use). Through a scripting language, the FMEBC supports model-to-model transformations, geometric restructuring, geometric and semantic filtering, datum and projection transformations, various quality control operations, clipping, overlay etc.

3.3 Published interfaces

Moving away from files, formats and encoding entirely, attention now turns to interoperability through interfaces. With such an approach, the

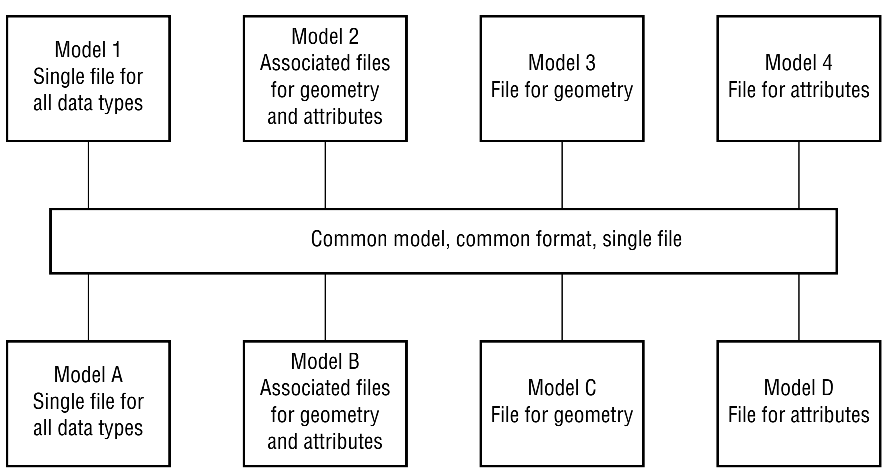

Fig 3. Data exchange through a common data model and format.

internal structure of the data is irrelevant; instead, the emphasis is on the behaviour presented by the interface. In other words, the interface is able to provide or accept data in response to a request. How it does so does not matter. There is no assumption that the data behind the interface must match the data provided to it or by it. Instead of making the specification of a proprietary format publicly available, vendors may choose to develop an API and to allow other developers to read and write their data through the API. An API can hide complexity, provide a number of capabilities supported by the application, and lessen the chances of inadvertently degrading database integrity. Moreover, it can be designed to respond to queries as part of a data access server. Translation software, as well as other applications, may interface to data holdings through an API (see also Worboys, Chapter 26).

Although APIs offer obvious advantages, they also have several disadvantages. They may be very dependent on particular technological environments, thus limiting their use in general interoperability scenarios. Because of marketing and engineering considerations, most APIs cannot operate independently of the main GIS application. Thus, if proprietary data files can be detached or exported from a GIS and shipped to another system, the receiver will not have access to the source vendor's API to help read the data. Only some GIS products have associated APIs, and typically, those that are available have not been created to any standard specification. Data received through different vendor APIs may conform to very different models, making integration of the data problematic.

An alternative approach to achieving interoperability is to develop an industry-wide common interface, based on distributed computing technologies (see also Worboys, Chapter 26). This is the primary goal of the Open GIS Consortium (OGC). Having coined the term OpenGIS, the OGC has as its mission 'the full integration of geospatial data and geoprocessing resources into mainstream computing and the widespread use of interoperable geoprocessing software and geospatial data products throughout the information infrastructure' (OGC 1996a). The OGC is developing an interface definition referred to as the OpenGIS Specification (formerly known as the Open Geodata Interoperability Specification). Interfaces compliant with this specification can be incorporated directly into new systems and built onto legacy systems.

Two significant components comprise the OpenGIS Specification: the Open Geodata Model (OGM) and the Services Architecture. The OGM is a collection of data types and methods, organised into a hierarchical class library. It is comprehensive, in that it embraces fundamental geospatial (and ultimately spatio-temporal) data types, including their geometric representation, spatial reference, and semantic content. The Services Architecture provides the mechanism by which individual objects and their associated interfaces may be assembled into complex queries, transformations, analytical functions, and presentation directives. It also enables the construction of catalogues that allow users to identify, evaluate, and interpret complex geographical information dispersed throughout a network.

Figure 4 shows a common interface for each of a number of data stores and applications. When a request is made by an application, it travels to the interface of a data store or another application. The interface returns the requested information as containers of information (encapsulated objects or components), in accordance with a distributed computing platform (DCP) specification. Such DCPs have been defined in Unix, Microsoft, and operating system independent environments. Examples include the Common Object Request Broker Architecture (CORBA) from the Object Management Group and the associated OpenDoc component model from the CI Labs consortium; Object Linking and Embedding (OLE), Common Object Model, Distributed Common Object Model and Active X from Microsoft; and Java and Java Beans from SunSoft. Consequently, the OpenGIS Specification must be sufficiently abstract to be able to be implemented in very different technical environments (OGC 1996b).

The common interface approach will allow disparate applications and technologies to interoperate as a single system. Ideally, applications can be built with both data and processing components coming from anywhere on a network. However, to define the common interface, to implement it effectively for day-to-day operational needs, and to gain widespread acceptance of it will take a long-term, concerted effort. Another potential disadvantage is that the common interface may not be capable of delivering objects required for a given application; for example, specific data types or subsets of the data may be of interest, but the interface specification may not have addressed these

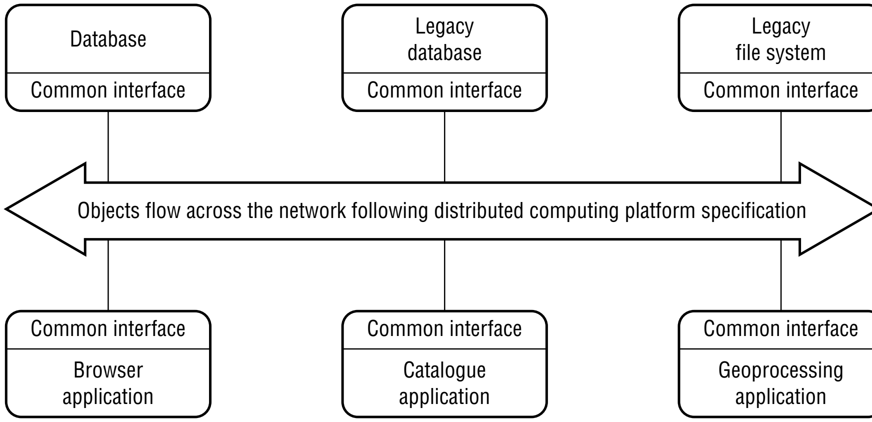

Fig 4. Interoperability through a common interface.

requirements adequately. We may see vendors supporting a common interface and extending it in line with their particular product capabilities. As has happened with SQL (section 3.4 below), the common interface once established may evolve over time to meet increasingly varied and complex needs.

3.4 Database developments

There has been a great deal of interest within the database world in managing atypical, non-tabular data, which together are considered as multimedia. The SQL language, long the backbone of the relational database market, was originally designed to handle only tabular data (Egenhofer and Kuhn, Chapter 28; Worboys, Chapter 26). However, the latest version, SQL3, is comparable to a full programming language and includes the ability to define abstract data types (ADT). An ADT may include functions; a polygon may have one function to return its boundary and another to return the number of holes within the polygon. An ADT is roughly equivalent to an object type in object-oriented technology. Instead of thinking about data in a database as consisting of rows of attributes with primitive domains (integers, real numbers, character strings, and the like), we can now include higher-level data types such as lines and polygons, each of which may have various functions as an inherent part of its definition.

The creation of a standardised set of data types, based on the ADT capability in SQL3, is the basis for the multimedia extensions known as SQL/MM and defined in the SQL3 language. One particular application area included in SQL/MM is Spatial (ISO/IEC JTC 1/SC 21 1996). These developments are significant because they enable databases to store and process a wide variety of data types, including spatial data types, all in the same environment. The enhancement of extended-relational and object-relational databases to geospatial applications parallels the migration of conventional APIs to more object-oriented designs encapsulating both data and methods. SQL/MM Spatial and the OpenGIS Specification are compatible standardisation efforts which will make it relatively easy for database vendors intending to comply with the former to also comply with the latter. Consequently, connectivity between different databases housing geospatial data and between such databases and other software products will become practical through the common interface defined by the OpenGIS Specification.

The different interoperability strategies described above have been developed to meet different technical objectives. Even so, there is a high degree of commonality in the underlying geospatial models. STDS, the first of these developments, strongly influenced aspects of DIGEST, and both influenced SAIF. SQL/MM Spatial and aspects of the OpenGIS Specification are related to SAIF, which in turn was modified to mesh with new ideas brought forward by these initiatives. Even with the diverse requirements driving all of these efforts, the underlying models have much in common. This is good news for organisations wishing to use various interoperability strategies (see Bédard, Chapter 29; Salgé, Chapter 50).

4 GEOGRAPHICAL DATA MODELLING

Data management and analysis demand that basic notions about data types be established through data models. With geographical data this implies that general-purpose concepts of geographical phenomena be available, as well as definitions of space and time. On the one hand, such models must be practical to understand and to implement; on the other, they must be sufficiently sophisticated and robust to describe data types applicable to a wide range of geographical applications. The modelling techniques used will influence the ease with which such concepts can be portrayed and implemented in different systems.

Geographical data modelling may make use of both text-based and graphical techniques. Examples of the former are INTERLIS, for data exchange of cadastral data (GEOHUB 1996), Class Syntax Notation, the data modelling language of SAIF (Geographic Data BC 1995), and SQL/MM Spatial (ISO 1996), designed specifically for database applications and under the auspices of the ISO. Graphical methods in common use include OMT (Rumbaugh et al 1991), the Unified Modelling Language (Booch et al 1996), Syntropy (Cook and Daniels 1994), and extended entity-relationship methods such as Designer/2000 (Oracle 1996). Bédard (Chapter 29) has also developed an approach with icon extensions to represent spatial and temporal constructs. Most of these techniques are object-oriented and provide a great deal of flexibility.

4.1 Abstraction of geographical phenomena

Fundamental to the interoperability problem is that different groups of users have different Earth models, which in turn manifest themselves in the representation of geographical information (Mark, Chapter 7; Martin, Chapter 6; Peuquet, Chapter 8; Raper, Chapter 5). Of significance is how perceived elements of the landscape are described and modelled, and how queries and analytical services on such data are defined and implemented. The most basic questions concern the conceptualisation of geographical entities and the space and time in which they reside. Such entities, existing in the real world, are represented in computer systems by units of data referred to as features or geographical objects, which in turn may be categorised into a general type of object termed Feature (using the OGC terminology, and equivalent to GeographicObject in SAIF).

4.1.1 Features

A *feature* is a representation of a phenomenon in object space, typically modelled through a series of attributes, including position. More specifically, it has a set of properties which includes its spatial representation, manifested through a geometric object and an associated spatial or spatio-temporal reference system. The set of properties may also include other properties, of a non-spatial nature. A number of these properties, in conjunction with the feature type, establish the semantics of the feature. A road segment may be defined as having various attributes including the number of lanes, the surface material, average traffic flows, update history, and a spatial representation as a geometric object such as an arc or curve. The coordinates of the points defining the geometry, and the implied geometric shape between the points, are meaningful in the context of a spatial reference system, including horizontal and vertical data and a planar or curvilinear projection.

A feature may also be composed recursively of other features, such as with a road network composed of roads or a farm consisting of agricultural fields and farm buildings. As described in the next section, some types of feature pertain to phenomena which vary continuously over a spatial extent, as with reflectance values in a satellite image, forest stand classification across a forest, or barometric pressure across a landscape.

Figure 5 provides a more formal definition of a feature (OGC 1966c), based on the OpenGIS Specification and using Syntropy. Every feature has an object identity which must be unique over space and time. It also has a property set, equivalent to a row in a relational database. The property set contains zero to many named properties (i.e. attributes). Each property may be simple (e.g. age represented as an integer) or complex (e.g. chemistry represented by a series of chemical attributes), and optionally may have constraints restricting its possible value. Zero to many properties are geometric; associated with a reference system, they specify the spatial or spatio-temporal extent of the feature. Geometric properties, which are types of properties, include coordinate geometry and may include other properties such as positional accuracy. A feature may be associated with other features, through various relationships, including spatial and temporal relationships. The association may also indicate containment, as previously noted. A user may define feature subtypes specific to his or her dataset or application.

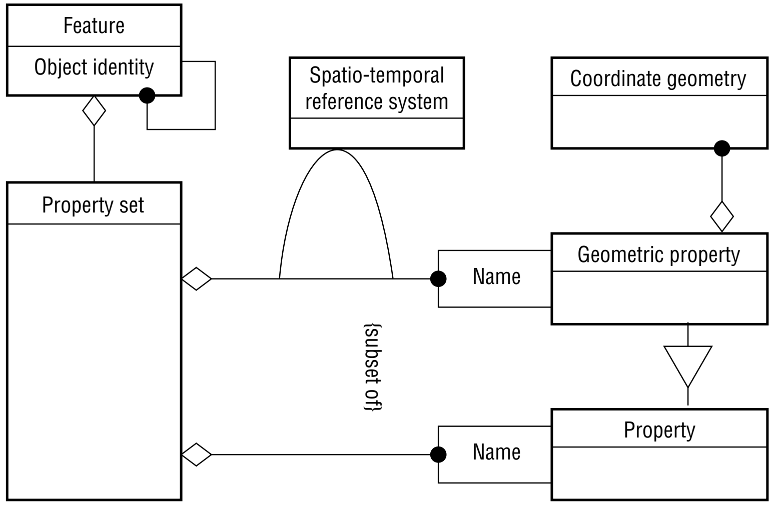

Fig 5. Feature, after the OpenGIS Specification.

4.1.2 Coverages

A point of ongoing discussion in the GIS literature has been the distinction between feature concepts based on discrete objects, and coverage concepts pertaining to continuous phenomena (Peuquet 1984). In the former case, the discrete object is generally treated as homogeneous with respect to its attribute values and usually the spatial boundaries of the object (e.g. the surface of a 3-dimensional object) are directly observable. Spatial location is often conceived as simply another attribute which distinguishes a geographical object from any other kind of object (Worboys, Chapter 26; 1995). Sometimes traditional attributes such as average elevation and slope are derived from the geometry and treated as additional dimensions in the multidimensional attribute space characterising typical analyses.

By contrast, a coverage (cf. the concept of a field discussed by Mark, Chapter 7) is a phenomenon in object space modelled as an attribute range (possibly complex) over some spatio-temporal domain. Coverage concepts imply a continuous surface in 2- or 3-dimensional space, with a given attribute value characterising any point on the surface. The concept of coverage may also be applied to solids (i.e. volumes), in 3-dimensional space. Mathematically, a coverage implies a computable function. Input to the function is any spatial or spatio-temporal position within a given spatial or spatio-temporal extent, and output is a value for a simple or complex attribute. Such a function may exist explicitly as a mathematical expression or it may simply be implied by the data type and structure.

Housing prices across a city could be represented as a trend surface through a polynomial or thin plate spline function. Breaks may exist on such surfaces, as with a cliff on an elevation map. With either prices or elevation, values may be displayed as contour lines or as a field of values at specific locations. A typical partitioned coverage of categorical thematic data shows bounded areas treated as internally homogeneous with discontinuities (i.e. boundaries) separating one area from another. Thus, a forest stand map or a map of countries can considered as defined by a discrete-valued function.

Error or uncertainty may be an intrinsic part of a phenomenon (Fisher, Chapter 13), as with the concentration of a given mineral calculated by a Kriging function, which produces both a predicted surface of the concentration and an associated error surface. These surfaces may be represented as interrelated coverages, or as a single coverage with a pair of predicted and error values as output. Geometrically, a simple surface may be expressed as contours, as a triangulated irregular network, as a set of points at locations of interest, etc. Fuzzy boundaries on categorical maps may imply fuzzy membership; land-use categories across the landscape may be represented by a set of membership functions, each giving the likelihood of belonging to a given category (Fisher, Chapter 13). In this case, the series of individual functions may be treated as a single function, with geographical position as input and a set of likelihood estimates as output.

From a modelling perspective the question is: what is the relationship between feature and coverage? Instead of treating them as two entirely separate types of objects, coverage may be modelled as a specialised type of feature which either explicitly or implicity stores a function able to provide an attribute value for any point across a given extent. Figure 6 presents a simplified view derived from the OpenGIS Specification (OGC 1996c). For reasons of clarity, the figure shows Geometry instead of Geometric Property, Coordinate Geometry, and Spatial Temporal Reference System (which are inherited from Feature).

Practically, a coverage and its associated function can take on one of several forms:

- a partitioned coverage with each area or patch defined by a discrete value;
- a set of values at specific point locations and with or without breaks or surface discontinuities;
- rasters of images and other similarly structured data;
- networked patches (e.g. triangulated irregular networks) with associated values;
- purely analytical functions;
- combinations of these such as a series of patches with an analytical routine pertaining to each or all patches.

From an interoperability perspective, the first four options should be reasonably easy to handle because they relate to well-known data structures. The fifth and sixth options pose the problem of representing the analytical functions and their programming interfaces in a universally acceptable way. Technically this requires that an interface to a dataset or application must not only be able to retrieve specific values, but also be able to specify parameters for an analytical method used at the server to generate the values. Alternatively, the methods may be returned to the requesting client, encapsulated with the raw data (see Goodchild and Longley, Chapter 40).

4.2 Space, time, and object identity

Concepts of space and time require (1) spatial and temporal reference systems and (2) geometric and temporal constructs which define position in the context of these systems (Couclelis, Chapter 2). The reference systems define the coordinate space, which typically has either two or three spatial dimensions and optionally a temporal dimension (Peuquet, Chapter 8).

A variety of types of spatial coordinate systems are in use around the world for different applications and in different areas. They include simple rectangular systems used in CAD environments to Earth-related systems defined through combinations of the reference ellipsoid, horizontal adjustment system, vertical adjustment system or surface, and projection. Because of extensive standardisation efforts in the geodetic realm, models of spatial reference systems are quite similar in most geoprocessing applications. A major exception has to do with linear referencing systems. Used extensively in road networks, they are based on driven distance from arbitrary reference points.

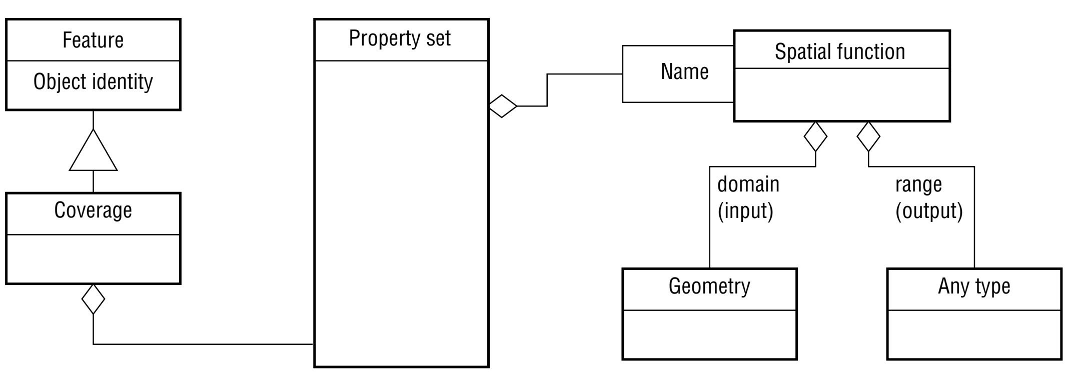

Fig 6. Coverage definition after OpenGIS Specification.

Interoperability with transportation and hydrological applications is complicated by the proliferation of linear referencing systems and discrepancies between positions defined in such systems and their correct geographical positions.

Time may be incorporated in two very different ways, either as part of a coordinate system, or as a time, a date, a time stamp (time and date together), or as an interval. The time may be specified with respect to universal time coordinates (UTC) or a local time zone. Less obvious issues of time are to distinguish among, and have access to, the actual time of an event, the time of observation, the time an object first entered a data store, the time the object was modified or replaced in a data store, and the time an object was retired or deleted from a data store (Peuquet, Chapter 8).

Closely tied to those notions of time is the issue of object identity. If a river changes course, we may consider that the geometry of a river segment has changed, or alternatively, we may state that the original segment has retired and been replaced by a new segment (Veregin, Chapter 12). Such concerns are of fundamental interest to interoperability if queries about changes in the landscape are to be answered. The problem is compounded by the requirements of different applications; it may be that a database must be able to present a number of views, depending upon the context of the query.

5 SHARING INFORMATION

5.1 Information communities

Information communities (Lake 1996) are groups of users, both data providers and consumers, who share digital data. To create a Geospatial Information Community or GIC (OGC 1996c) typically requires acceptance of particular feature (and coverage) models, reference systems, and geometries. Where models from user to user are divergent, the users may still form a community, so long as they have the means to exchange information through an interoperability strategy.

A number of organisations have created information communities around particular needs. Three of these are briefly reviewed. The Petrotechnical Open Software Corporation (POSC 1995) is a not-for-profit organisation funded by the oil industry with the objective of establishing and promoting standards for information sharing specifically within the exploration and production sectors of that industry. Part 4 of DIGEST (see section 3.2) defines the Feature and Attribute Coding Catalogue (FACC; Digital Geographic Information Working Group 1995b). The FACC provides standardised definitions and codes for features in ten categories (Culture, Hydrography, Hypsography. Physiography, Vegetation, Demarcation, Aeronautical Information, Cadastral, Special use, and General). It is the intention of the Digital Geographic Information Working Group that all NATO defence forces adopt these definitions and use them with the DIGEST transfer format. Originally developed by private industry and now under the auspices of the Comité Européen de Normalisation, TC 278, the Geographic Data Files standard (GDF: European Commission DGXIII 1996) includes models of road network features (Salgé, Chapter 50). GDF is playing a prominent role, especially in Europe, in the development of intelligent transportation systems.

Efforts are under way to create general-purpose geospatial information communities. Under the Federal Geographic Data Committee, the National Spatial Data Infrastructure (FGDC 1996) is a long-term effort to share geospatial data throughout the USA. It has been structured around two major activities: the Clearinghouse, which is a metadata cataloguing and searching model for enabling access to distributed geodata resources, and the Framework, which is focused specifically on characterising a collection of datasets that are of national interest and whose feature attributes may be standardised (Guptill, Chapter 49).

5.2 Schemas, schema fusion, and schema transformation

A set of data type definitions, i.e. data models, constitutes a *schema*. It provides a systematic and coherent description of the content and organisation of a dataset or collection. It can include the semantic representation of features, the details of attribute names and types, the dictionary defining the structure of geographical objects and their spatial and non-spatial attributes, metadata providing a synoptic view of geographical information, and a thesaurus of related terms. A government agency may define schemas pertaining to the types of data they collect and distribute, and similarly, members of an information community may define a schema meaningful to them.

Key interoperability issues include how schemas are represented, how a given user-defined schema relates to standardised geospatial concepts (e.g. feature, reference system, geometry), and how different user-defined schemas relate to one another. Geospatial standards provide a high degree of semantic consistency insofar as fundamental geospatial concepts are concerned. Outside these concepts, however, semantic consistency remains an issue. *Schema fusion* involves creation of a common schema capable of describing all or most information of interest. At times a more practical alternative is to establish a series of model-to-model transformations, such that a given schema is redefined in a way most appropriate to the recipient. This latter approach is more flexible but only possible with a detailed knowledge of the respective schemas and *schema transformation* software.

5.3 Metadata and catalogues

Metadata, or data about data, may be used to give a high-level description of a data collection, including such items as the feature types, ownership, quality, lineage, spatial referencing, currency, and version (Guptill, Chapter 49). Metadata are intended to serve many objectives, ranging from dataset documentation, to provision of concise definitions of the dataset's structure and organisation, to a basis for browsing and searching (Goodchild and Longley, Chapter 40). One of the more influential developments is the Content Standards for Digital Geospatial Metadata from the US Federal Geographic Data Committee (FGDC 1994). The standard specifies the metadata content only; it does not state how such data should be stored or accessed. It also does not require that schemas or detailed data dictionaries be available.

A *catalogue* is a collection of metadata descriptions which may apply to various levels of aggregation or granularity. For example, it may pertain to collections of datasets (such as all USGS 7.5' quadrangle maps), to a single dataset (all data on a given map), or to given types of data only (all roads on the entire series or on some subset of it). Acting as a gazetteer or index, a catalogue serves as an authoritative reference on one or more data collections and may include a data dictionary and formal data models. Potential data users may browse the catalogue to determine the data's relevance, extent, cost, etc., as well as valid ways to query the data. The catalogue may even control access through some sort of authentication process. A significant concern with catalogues is to ensure that the metadata they contain is current. If the catalogue includes descriptions of a large number of active datasets, maintenance requires ongoing cooperation from the data providers. Such cooperation is particularly practical if the data providers are all part of an information community.

6 CONCLUSIONS

The GIS interoperability strategies described in this chapter are not mutually exclusive. Data managers and GIS practitioners may use them in complementary ways, as appropriate for specific environments or data holdings. Initially vendors are likely to develop interoperability among their own products. As pressure from specific communities grows and as the underlying technology matures, both vendors and third parties can be expected to offer true, multi-vendor interoperability solutions. It is easy to imagine purchasing criteria which include meeting specific interoperability capabilities.

GIS interoperability developments are significant for another reason as well: they are contributing to the migration away from the monolithic systems which have dominated the GIS market for so long. Databases, browsers, smart translators, and geoprocessing tools are now being coupled through adherence to open standards and specifications. The backbone of much GIS activity is likely to be network-accessible databases, with connectivity requirements met by common interfaces. The big winners will be information communities gearing up to take advantage of these new opportunities.

References

Booch G, Rumbaugh J, Jacobson I 1996 *The unified modelling language for object-oriented development. Documentation set Version 0.9 addendum*. Santa Clara. *http://www.rational.com/ot/uml.html*

CEN/TC 287 Secretariat 1996 *CEN/TC 287 Geographic Information. http://www.statkart.no/sk/standard/cen/*

Cook S, Daniels J 1994 *Designing object systems: object-oriented modelling with syntropy*. Englewood Cliffs, Prentice-Hall

Digital Geographic Information Working Group 1995a *Digital geographic information exchange standard*. Version 1.2a. *http://132.156.33.161/Engineer/DIGEST_1.2a/cover.htm*

Digital Geographic Information Working Group 1995b *Feature attribute and coding catalogue.* Version 1.2a. *http://132.156.33.161/Engineer/DIGEST_1.2a/covers/part4.htm*

European Commission DGXIII 1996 GDF home page. *http://205.139.128.5/ehq/gdf*

FGDC 1994 *Content standards for digital geospatial metadata,* Washington DC, Federal Geographic Data Committee. *http://www.fgdc.gov/metadata*

FGDC 1996 *National Spatial Data Infrastructure,* Washington DC, Federal Geographic Data Committee. *http://www.fgdc.gov/nsdi2.html*

Geographic Data BC 1995 *Spatial archive and interchange format: formal definition.* Release 3.2. *http://www.env.gov.bc.ca/gdbc/saif32*

Geographic Data BC 1996 *Feature manipulation engine BC. http://www.env.gov.bc.ca/gdbc/fmebc*

GEOHUB 1996 *INTERLIS. http://www.geo.unizh.ch/~keller/geohub.html#tagOO2*

ISO/IEC JTC 1/SC 21 1996 *SQL multimedia and application packages (SQL/MM) – Part 3: spatial*. ISO/IEC JTC 1/SC 21 N 10441, ISO/IEC CD 13249-3:199x (E) SQL/MM:MAD005. *ftp://speckle.ncsl.nist.gov/isowg3/sqlmm/MADdocs/*

ISO/TC 211 Secretariat 1996 *ISO/TC 211 Geographic Information/Geomatic. http://www.statkart.no/isotc211/*

Lake R W 1996 Information communities support information sharing. *GIS World* 9(2): 72–3

Litwin W, Mark L, Roussopoulos N 1990 Interoperability of multiple autonomous databases. *ACM Computing Surveys* 22: 265–93

OGC (Open GIS Consortium) 1996a *Vision and mission statements. http://www.opengis.org/vision.html*

OGC (Open GIS Consortium) 1996b *The OpenGIS guide, introduction to interoperable geoprocessing, Part 1. http://ogis.org/guide/guide1.html*

OGC (Open GIS Consortium) 1996c *The OpenGIS abstract specification*. *http://www.opengis.org/public/abstract.html*

Oracle Corporation 1996 *Oracle Designer/2000*. Release 1.3

Petrotechnical Open Software Corporation (POSC) 1995 *POSC technical information*. *http://www.posc.org/technical_toc.html*

Peuquet D J 1984 A conceptual framework and comparison of spatial data models. *Cartographica*. 21: 66–133

Rumbaugh J, Blaha M, Premerlani W, Eddy F, Lorensen W 1991 *Object-oriented modelling and design*. Englewood Cliffs, Prentice-Hall

US Geological Survey 1996 *Spatial data transfer standard. http://mcmcweb.er.usgs.gov/sdts/standard.html*

UCGIS (University Consortium for Geographic Information Science) 1996 *The UCGIS research agenda: interoperability*. *http://www.ucgis.org*

Worboys M F 1995 *GIS, a computing perspective*. London, Taylor and Francis