

Chapter 12

TerraLib: An Open Source GIS Library for Large-Scale Environmental and Socio-Economic Applications

Gilberto Câmara, Lúbia Vinhas, Karine Reis Ferreira, Gilberto Ribeiro de Queiroz, Ricardo Cartaxo Modesto de Souza, Antônio Miguel Vieira Monteiro, Marcelo Tílio de Carvalho, Marco Antonio Casanova and Ubirajara Moura de Freitas

Abstract This chapter describes TerraLib, an open source GIS software library. The design goal for TerraLib is to support large-scale applications using socio-economic and environmental data. TerraLib supports coding of geographical applications using spatial databases, and stores data in different database management systems including MySQL and PostgreSQL. Its vector data model is upwards compliant with Open Geospatial Consortium (OGC) standards. It handles spatio-temporal data

Gilberto Câmara

National Institute for Space Research (INPE), Av dos Astronautas 1758, 12227-010, São José dos Campos, Brazil, e-mail: gilberto@dpi.inpe.br

Lúbia Vinhas

National Institute for Space Research (INPE), Av dos Astronautas 1758, 12227-010, São José dos Campos, Brazil, e-mail: lubia@dpi.inpe.br

Karine Reis Ferreira

National Institute for Space Research (INPE), Av dos Astronautas 1758, 12227-010, São José dos Campos, Brazil, e-mail: karine@dpi.inpe.br

Gilberto Ribeiro de Queiroz

National Institute for Space Research (INPE), Av dos Astronautas 1758, 12227-010, São José dos Campos, Brazil, e-mail: gribeiro@dpi.inpe.br

Ricardo Cartaxo Modesto de Souza

National Institute for Space Research (INPE), Av dos Astronautas 1758, 12227-010, São José dos Campos, Brazil, e-mail: cartaxo@dpi.inpe.br

Antônio Miguel Vieira Monteiro

National Institute for Space Research (INPE), Av dos Astronautas 1758, 12227-010, São José dos Campos, Brazil, e-mail: miguel@dpi.inpe.br

Marcelo Tílio de Carvalho

Catholic University of Rio de Janeiro (PUC-RIO), Rua Marquês de São Vicente, 22522. 453-900 Rio de Janeiro/RJ, Brazil, e-mail: tilio@tecgraf.puc-rio.br

Marco Antonio Casanova

Catholic University of Rio de Janeiro (PUC-RIO), Rua Marquês de São Vicente, 22522. 453-900 Rio de Janeiro/RJ, Brazil, e-mail: casanova@tecgraf.puc-rio.br

Ubirajara Moura de Freitas

Space Research and Applications Foundation (FUNCATE), Av. Dr. João Guilhermino, 429 – 18th floor 12210-131 São José dos Campos, SP, Brazil, e-mail: bira@geo.funcate.org.br

with Open Geospatial Consortium (OGC) standards. It handles spatio-temporal data types (events, moving objects, cell spaces, modifiable objects) and allows spatial, temporal, and attribute queries on the database. TerraLib supports dynamic modeling in generalized cell spaces, has a direct runtime link with the R programming language for statistical analysis, and handles large image data sets. The library is developed in C++, and has programming interfaces in Java and Visual Basic. Using TerraLib, the Brazilian National Institute for Space Research (INPE) developed the TerraView open source GIS, which provides functions for data conversion, display, exploratory spatial data analysis, and spatial and non-spatial queries. Another noteworthy application is TerraAmazon, Brazil's national database for monitoring deforestation in the Amazon rainforest, which manages more than 2 million complex polygons and 60 gigabytes of remote sensing images.

12.1 Introduction

Recent advances in spatial databases have changed both the nature and process of geographic information system (GIS) software development. Spatially-enabled database management systems (DBMS) such as PostgreSQL empower a transition from monolithic GIS with hundreds of functions to a generation of spatial information applications tailored to suit specific user needs. These capacities have been a major boon for the free and open source geospatial (FOSS4G) community, many members of which are using the new generation of databases to build unique and innovative applications.

One of the expected impacts of open source software (OSS) is its benefits for developing nations. As Weber (2004) points out, combining OSS with the technical workforce available in developing countries can enable technology transfer. He states, "Of course information technology and open source in particular is not a silver bullet for long-standing development issues; nothing is. But the transformative potential of computing does create new opportunities to make progress on development problems that have been intransigent" (Weber 2004 p. 254).

Following from this point, GIS is a key technology for developing nations in domains such as environmental protection, urban management, agricultural production, deforestation mapping, public health assessment, crime-fighting, and socio-economic measurements. However, the demands of these applications go well beyond the current specifications of the Open Geospatial Consortium (OGC). Large-scale environmental and socio-economic applications compel FOSS4G to include significant spatial analysis capacities to meet the needs of end-users (Goodchild 2003). Hence, FOSS4G should incorporate research advances in areas such as spatio-temporal data models (Erwig and Schneider 2002; Hornsby and Egenhofer 2000), geographical ontologies (Fonseca et al. 2002), spatial statistics and spatial econometrics (Anselin 1999), cellular automata (Couclelis 1997), and environmental modeling (Burrough 1998). These topics have largely been outside the reach of the GIS user community due to a general lack of widely available tools that support

them. Incorporation of some of these new techniques into GIS applications is necessary for the user community to extract the full potential of spatial databases.

With this motivation, TerraLib was developed as an open source GIS software library that extends object-relational DBMS technology to support spatio-temporal models, spatial analysis, spatial data mining, and image databases. The design goal for TerraLib is to support large-scale applications using cadastral, socio-economic and environmental data. This goal was a mandate of the main organization that supports TerraLib, the Brazilian National Institute for Space Research (INPE). INPE is Brazil's primary institution for space science and technology. Its mission includes building satellites, developing environmental applications, and producing weather and climate forecasts. Since 1984, INPE has had a research and development division for GIS to support its actions in earth observation and to promote GIS and remote sensing technology in Brazil. The two other main project partners are the Computer Graphics Group (TecGraf) of the Catholic University of Rio de Janeiro (PUC-RIO) and FUNCATE, a non-profit foundation that develops GIS applications using OSS. All organizations involved in the TerraLib project share the same general design goals. Thus, TerraLib is a project with long-term support and a stable and secure working environment for its developers. This chapter describes the TerraLib library, explains the main design decisions, and points out how the library incorporates research results from GIScience in its development.

12.2 Challenges for Innovation in FOSS4G

The OGC specifications noted above provide a sound basis for developing FOSS4G projects. However, many applications need tools which go beyond these specifications. Thus, one of the lines of growth in FOSS4G is to provide new tools for application developers. However, there are pitfalls. Building innovation in open source GIS is a threefold challenge. Given the design goals for the project discussed in this chapter, the first step required selecting, from the large body of GIScience literature, those advances that are relevant to the project's objectives. These advances then need to be implemented in industrial-strength code. The final hurdle is documenting these features and sharing them with the broader FOSS4G development community.

A basic design objective for TerraLib was to support innovative applications to help people and protect the environment. Thus, current GIS research was first evaluated, and ideas and proposals were selected that were relevant to the design goals. This led to concentration in the following three areas:

- (a) *Spatial Statistics*: since Anselin's pioneering work on spatial analysis (Anselin 1989), promising advances have appeared in the field of spatial statistics and spatial data mining (Anselin 1995; Fotheringham et al. 2002; Openshaw and Albanides 2001; Martin 2003). The main focus of these contributions is to improve the ability to extract information for socio-economic data. This is relevant to public policy applications of GIS.

- (b) *Spatio-temporal Models*: there are two broad categories of spatio-temporal objects. The first concerns *moving objects*. Moving objects relate to, for example, information about spatial and temporal positions of planes, storms or automobiles. The widespread relevance of location-based applications has motivated developments in the field of moving object databases (Güting and Schneider 2005). There is a large research area in algorithms and query methods for moving objects (Sistla et al. 1997). The second type concerns *evolving objects* that do not move, but whose geometry, topology and properties change. They arise when changes that occur in, for example, cadastral GIS or in land cover patterns are considered (Medak 2001). Evolving objects are important for environmental models, which depict the temporal evolution of a pattern in a landscape. Examples of environmental models include land change models, epidemiological studies, population flows, and ecological mapping (Burrough 1998; Veldkamp and Fresco 1996).
- (c) *Remote Sensing, Image Processing, and Image Databases*: remote sensing satellites are the most significant source of new data about our planet, and remote sensing image databases are the fastest growing archives of spatial information. New high resolution optical sensors and polarimetric radars have improved application areas such as environmental monitoring and urban management. There are important recent advances in object-oriented segmentation and classification, and in remote sensing data mining (Blaschke and Hay 2001; Navulur 2006; Aksoy et al. 2004). It is also important to include support for raster data handling in open-source DBMS, following the research results of Chang et al. (1988) and DeWitt et al. (1994).

To translate these ideas to industrial-strength code, the developers of TerraLib first undertook various research projects and published the results from these (Pedrosa et al. 2002; Almeida et al. 2003; Vinhas et al. 2003; Ferreira et al. 2005; Silva et al. 2005; Assunção et al. 2006; Feitosa et al. 2007). These results enabled the TerraLib development team to assess the potential benefits of each technique, as well as the trade-offs needed to generate production code. Software engineering tools for GIS were also examined during this process. One of the conclusions from this was to confirm the usefulness of *design patterns* and *generic programming* as a basis for achieving reuse in GIS software development (Câmara et al. 2001; Vinhas et al. 2002).

The last and most difficult problem is sharing the resulting code with the FOSS4G community. Many of the new tools and techniques might be unfamiliar for FOSS4G developers and practitioners. Hence, there is a need to explain not only the code, but also the ideas behind it. Experience has shown that face-to-face workshops and meetings are the best way to discuss new ideas and their implementation. A second-best alternative is writing detailed documentation, which is not easy to achieve in open source projects (see Chap. 2). Developers have to work hard to share their results, and the TerraLib team is aware of this challenge.

12.3 The Design of TerraLib

This section discusses the requirements and design rationale for TerraLib. It presents the alternatives considered at various points, and explains the final choices that were made. The discussion explains how product requirements led to the software architecture, the conceptual model, and extensions to the basic OGC specifications.

12.3.1 Product Requirements

The main goal for TerraLib led to the following needs:

- (a) *Ease of customisation*: developers should require little effort to use the library to develop their applications. They should concentrate only on specific user needs, and the library should provide powerful abstractions that cover the common needs of a GIS application.
- (b) *Upward compatibility to the OGC simple feature data model*: considering the impact and popularity of the OGC specifications, a TerraLib spatial database should be compatible with the OGC simple feature specification (SFS). This was not an original project requirement. When the project started in 2002, the developers initially underestimated the impact and extent of the OGC specifications. Hence, TerraLib's code was redesigned (from Version 3.2 to Version 4.0) to satisfy the need for conformance.
- (c) *Decoupling applications from the DBMS*: the library should handle different object-relational databases transparently.
- (d) *Support for large-scale applications*: to be useful for environmental and socio-economic applications, the library should provide efficient storage and retrieval of hundreds of thousands of spatial objects.
- (e) *Extensibility*: a GIS library should be extensible and accessible by other programmers. Introducing new algorithms and tools should not affect already-existing code.
- (f) *Enabling spatio-temporal applications*: emerging GIS applications need support for different types of spatio-temporal data, including events, mobile objects, and evolving regions.
- (g) *Remote sensing image processing and storage*: the library should be able to handle large image databases, and inclusion of image processing algorithms should be easy.
- (h) *Spatial analysis*: there should be support for spatial statistical methods to improve the ability to extract information from socio-economic data.
- (i) *Environmental modelling*: there should be support for environmental and urban models, including dynamic models using cellular automata.

To respond to issues (a) and (b), TerraLib has a strong conceptual model, as explained in Sect. 12.3.2. Points (c), (d) and (e) led to a software architecture described in Sect. 12.3.3. The last four issues are considered in Sects. 12.3.4–12.3.7.

12.3.2 Conceptual Model

This section describes TerraLib's conceptual model that was designed to support requirements (a) and (b) noted above. When designing TerraLib, the developers had to make numerous choices which are typical of software library design in general (Meyer 1990; Krueger 1992; Fowler et al. 1995). Apart from basic principles such as applicability, efficiency, ease of use, and ease of maintenance, there are important trade-offs. In this regard, consider two opposing visions:

- *Vision 1*: Libraries should take a minimalist approach. They should provide only primitive building blocks and include generators that can combine these blocks to yield complex custom applications. They should be split into independent modules, with as few dependencies as possible. The developer's focus can be narrowed to those modules that are of interest (Batory et al. 1993).
- *Vision 2*: Libraries should have strong ideas behind them. All the functionalities and modules should work well together. The idea is to maximize reuse by minimizing cognitive distance, which Krueger (1992, p. 136) defines as: "the amount of intellectual effort expended by software developers to take a software system from one stage to another". In this vision, the intellectual effort that software developers need to take a library and development of applications should be minimal. Application programmers use higher-level abstractions to build applications, and do not need to understand the details of the library's source code.

The choice between the two visions depends on a library's initial design goals. For libraries designed to be part of a larger software project, the first vision is the usual choice. Examples include libraries such as the Standard Template Library (STL) in C++ (Austern 1998), or the *shapelib* utility for GIS (Warmerdam 2007 – see Chap. 5). At the other extreme, libraries are designed to be easily extendible to build complete applications. One example includes libraries that use the model-view-controller (MVC) pattern (Krasner and Pope 1988) such as Java Swing (Elliott Eckstein et al. 2002). Libraries with strong concepts dictate how the user should develop the application.

TerraLib follows the second vision, since it aims to make it easy for programmers to develop end-user applications. To do this, the library needs to consider the semantic mismatch between relational databases and object-oriented applications. Relational databases store information in *tuples*, but GIS applications manipulate *objects*. A typical GIS application consists of four steps: (a) query the spatial database; (b) convert the query results (tuples) into objects; (c) manipulate these objects to create new objects; (d) display the resulting objects.

Thus, applications need to distinguish between data sources (the *spatial database*) and data targets (*the set of objects that must be manipulated and displayed*). To reduce the cognitive distance from OSS code to a deliverable application, the GIS developer needs a library that provides abstractions both for the data sources and for the data targets. These abstractions should support the four basic GIS components (query, conversion, manipulation, and display). Consequently, TerraLib supports the following abstractions:

- *Database*: a repository of information that contains data and metadata.
- *Layer*: a container of spatial objects that share a common set of attributes. Examples of layers are thematic maps (soil or vegetation maps), cadastral maps (map of land parcels in a city), or raster data such as satellite imagery. A layer knows its cartographic projection. Layers are inserted in the database by importing data from files or other databases, or by processing other layers. A layer stores the temporal evolution of the objects it contains.
- *Representation*: the geometric parts of data contained in a layer. TerraLib supports different geometries, including two-dimensional (2D) vectors (points, lines or areas), cell spaces, networks, triangulated irregular networks (TINs), and multi-dimensional rasters. The same data can have different representations (for example a city can be represented by the polygon that describes its political boundaries or by a point that represents its geometric centre).
- *Theme*: a theme contains a subset of the objects of a layer, produced by a selection. The selection may use attribute, spatial, or temporal conditions. Each theme has a set of presentation attributes for graphical display.
- *View*: this is a set of themes that are visualized or processed together. It defines a particular user's view of the database. A view has a unique cartographic projection, and the themes it contains are converted to this projection.
- *Visual*: this comprises a set of presentation attributes. Each theme has a unique "visual". A visual includes choroplethic filling and contour colours for polygons, thickness and colours for lines, or symbols for points.

TerraLib distinguishes between data sources and data targets. The abstractions of *database*, *layer*, and *representation* relate to the source domain and describe data organization and hierarchy. The ideas of *theme*, *view*, and *visual* relate to the target domain and describe data retrieval and presentation. A query in TerraLib retrieves tuples from layers, converts these tuples into a set of objects, and groups objects of the same type in themes. Thus, *layer* and *theme* are complementary abstractions. *Layers* organize spatial data in the database. *Themes* organize objects for manipulation and display. Similarly, *databases* and *views* are complementary concepts. A *database* organizes *layers* of spatial data. A *view* organizes *themes* containing spatial objects.

These concepts provide a set of higher-level abstractions on top of the OGC SFS, which are not part of the current OGC model. TerraLib stores these entities in a set of metadata tables, built when creating a new database. These metadata tables are kept updated as long as TerraLib manages the database. Should an OGC-compliant application access a TerraLib database, it will only access the tables described in the OGC model.

12.3.3 Software Architecture

This section discusses how TerraLib responds to requirements (c), (d) and (e) as stated in Sect. 12.3.1 (*DBMS-independence*, *efficiency*, and *extensibility*). To address

these issues, it was decided to use the C++ programming language for development. The developers had previous experience and had developed many algorithms in C++ as part of SPRING, their earlier GIS project (Câmara et al. 1996). Existing DBMS such as PostgreSQL provide native interfaces in C++. Also, C++ helps with the use of generic programming (Alexandrescu 2001) and design patterns (Gamma et al. 1995).

The developers chose an architectural design that has a kernel and a periphery. Maintenance of the kernel is the responsibility of a core team composed of a few senior programmers. Other contributors use the library's core to add new algorithms that test the library's core for extensibility and robustness. This follows the approach used for successful OSS products such as Linux, PostgreSQL, and Apache, which all have a kernel whose maintenance is the responsibility of a small team. Contributions from the community occur at the external layers. As an example, out of more than 400 developers, the top 15 programmers of the Apache Web server contribute 88% of added lines (Mockus et al. 2002). TerraLib's architecture has four parts, as shown in Fig. 12.1:

- *Kernel*: the core of TerraLib provides a set of spatio-temporal data types, code for cartographic projections and topological spatial operators, an API for storage and retrieval of spatio-temporal objects in databases, and classes for controlling visualization of spatial data.
- *Drivers*: modules that specialize the kernel's generic database application programming interface (API) to allow access to DBMS such as PostgreSQL (with or without the PostGIS spatial extension) or MySQL, and to external files in both

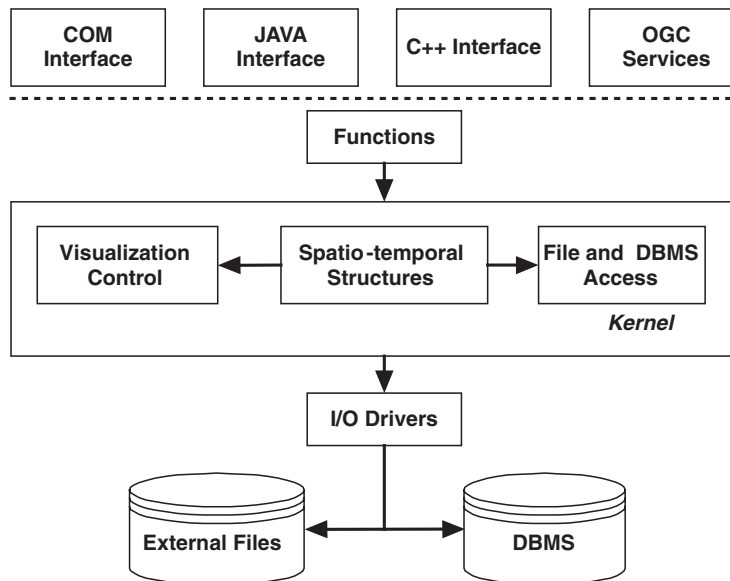


Fig. 12.1 TerraLib software architecture

open and proprietary formats. Basic maintenance and upgrade is the responsibility of the project core team.

- *Functions*: algorithms that use the kernel API. Typical functions include spatial analysis, query, and simulation languages. The functions are designed to allow external contributions.
- *Interfaces*: different interfaces to the TerraLib library that allow software development in different environments (Java, COM) and the support for OGC services such as Web Map Services (WMS), Web feature services (WFS) and Web coverage services (WCS).

The core of TerraLib's kernel is its set of spatio-temporal data types and its methods for query processing. The OGC Simple Feature Geometry model is used for storing basic vector geometries. Extra metadata tables support the abstractions described in Sect. 12.3.2. For a full description of the metadata, see the library's documentation (Vinhas et al. 2007). One key need is efficiency. The developers have spent much effort on issues such as indexing techniques and computational geometry algorithms (Queiroz 2003; Rodrigues et al. 2005; Rodrigues et al. 2006). With this work, the library now supports large-scale geographical databases, as discussed in Sect. 12.4.

The I/O drivers provide the interface between the kernel and the various DBMS and file formats. The library handles different object-relational databases, using a generic database API that handles the specific features of each DBMS. Using this API, a TerraLib programmer can work at an abstract level. TerraLib hides the differences between products such as PostgreSQL/PostGIS and MySQL from the programmer (Ferreira et al. 2002).

As noted earlier, the design of TerraLib *functions* aims for extensibility, as introducing new algorithms and tools should not affect existing code. Adoption of the principles of generic programming and design patterns helped achieve extensibility in this regard (Câmara et al. 2001). Three design patterns were found to be especially useful:

- *Factory*: this pattern provides an interface for creating an object, but lets subclasses decide which class to instantiate. In GIS, it is useful to include new functions without changing existing code. For example, there are hundreds of cartographic projections. When code for a new projection is inserted in TerraLib, it tells the projection factory about its existence and the projection factory calls this new code when it is needed.
- *Strategy*: provides an interface to a family of algorithms, and makes them interchangeable. In GIS, the strategy pattern is useful when there are different ways of performing the same function. This occurs often in image processing. For example, there are many different types of image filters. By using the strategy pattern, a programmer can use the same code for different filters.
- *Iterators*: TerraLib uses iterators to decouple algorithms from data structures. For example, to compute a histogram it is not essential to know if the data are a set of points, a set of polygons, a grid, or an image. The algorithm only needs to look into a list and get the values of the items that satisfy a certain property (for

example, those that are closer in space than a specified distance). In a similar way, most spatial analysis algorithms can be independent of spatial data structures and described only by their properties (Vinhas et al. 2002).

12.3.4 Raster Data Handling

As noted earlier, TerraLib handles raster data types as well as vector data. All raster data types are handled in a unified way, using an interface with two main methods, namely one to set the value of a point on a multidimensional raster and one to recover its value. The library provides decoders for different raster data formats, and iterators for accessing image data and developing image processing algorithms. The decoders are responsible for handling the particularities of each data source. The iterators are specialized pointers that traverse a raster in a predefined way (for example, only inside a given polygon). They hide the internal details of the raster data (Vinhas et al. 2003).

The library has drivers for storing raster data in different DBMS. TerraLib uses indexing and compression to achieve good performance in a standard DBMS, even for large satellite images. Indexing combines tiling and multi-resolution pyramids. Multi-resolution pyramids store the raster data at various sizes and degrees of resolution. Each resolution level is divided into tiles. A tile has a unique bounding box and a unique spatial resolution level, which are used to index it. Figure 12.2 shows a pictorial representation of the tiling and multi-resolution storage model.

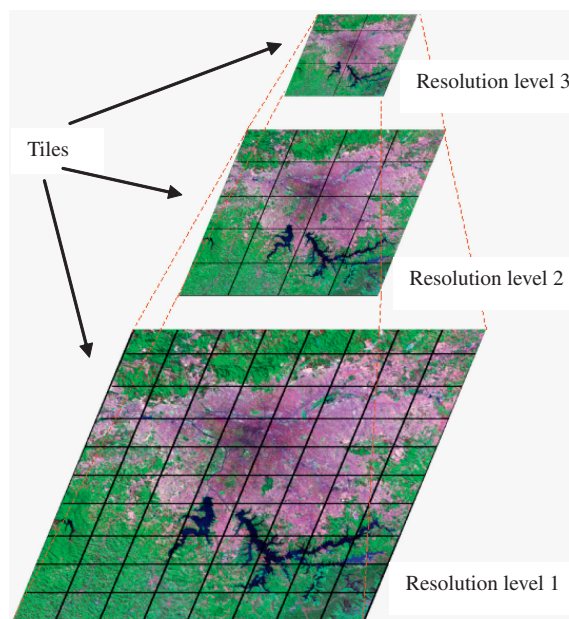


Fig. 12.2 TerraLib raster data indexing

The multi-resolution pyramid approach is useful for display of large data sets, avoiding unnecessary data access. TerraLib stores the whole pyramid. To compensate for the extra storage needs, it applies lossless compression to the individual tiles. When retrieving a section of the data, only the relevant tiles are accessed and decompressed.

TerraLib provides a large set of image processing algorithms including filters, segmentation, classification, mixture models, and geometric transformations. The image processing algorithms have a common interface for receiving as input instances of the raster API and a set of parameters. Two particularly important algorithms are the object-oriented segmentation and region classifier algorithms developed by INPE, originally as part of the SPRING software (Cámara et al. 1996). These algorithms were extensively validated for extracting land use patterns in tropical forests, and were favourably reviewed in a recent survey (Meinel and Neubert 2004).

Creating new image processing algorithms is straightforward. The library has a set of standard protocols that combine the *Factory* and *Strategy* design patterns. First, a programmer develops a new algorithm (for example, a filter), then instructs TerraLib to use the proposed strategy to filter the image. The programming manual provides further details as to the subsequent steps that need to be followed to integrate the new code (Vinhas et al. 2007).

12.3.5 Spatio-temporal Queries

There are numerous different ways to record spatio-temporal information in a database. The main alternatives are to (a) provide snapshots of data, (b) store sequences that describe the temporal evolution of spatial objects, or (c) store both objects and events that change them (Hornsby and Egenhofer 2000; Grenon and Smith 2003; Galton 2004; Worboys 2005). TerraLib adopts a dual perspective, namely archiving fields (stored in raster data) as snapshots and objects (stored in vector data) as sequences. A spatio-temporal object in TerraLib is a sequence of static objects with the same identifier. Each static object is valid for one interval.

Storage and retrieval of spatio-temporal objects in a DBMS needs more abstractions beyond those discussed in Sect. 12.3.2. TerraLib distinguishes four types of data stored in layers, namely *static* (unchanging data), *events* (singular occurrences such as crime events), *moving objects* (such as cars on highways) and *evolving objects* (such as cities, whose boundaries and attribute values change in time). All spatio-temporal objects that share the same attributes are converted to tuples of a layer (including timestamps) and stored together in a database. Metadata tables store information about different types of layers and identify which attributes of a layer store the timestamps associated to the temporal intervals. For example, consider an urban cadastre where all land parcels in a city for all intervals are stored together in a single layer. Grouping all objects together in this manner simplifies data handling. Inside the database, TerraLib uses optimization techniques for dealing with large data volumes.

As discussed in Sect. 12.3.2, a GIS application needs to transform database tuples (data sources) into objects that can be manipulated and displayed (data targets). In a purely static and non-temporal GIS, it is enough to use *themes* to group objects of the same type resulting from queries. In this case, all objects contained in a theme belong to the same interval.

When a *theme* of spatio-temporal data is retrieved from the database, it contains all spatio-temporal objects for the whole period when data are available. However, not all objects exist in all instances. Consider the case of real estate in an urban cadastre. Extracting data from the “parcels” *layer* will produce a *theme* composed of all parcels that ever existed in the cadastre. An application may need to use only those parcels that currently exist. In this case, a spatio-temporal selection is required inside a *theme*. To do this, TerraLib uses an extra concept referred to as a spatio-temporal object (*STObject*). An *STObject* is an individual entity that preserves its identity, but may change its location and the values of its attributes.

TerraLib provides a query processor to extract *STObjects* from themes (Ferreira et al. 2005, see Fig. 12.3). The query processor has a generic API for programmers, hiding data storage details. Algorithms can handle spatio-temporal data using only *STObjects* returned by the query processor. To perform a query, a programmer defines three different restrictions (spatial, temporal, and attribute). The spatial restrictions use the OGC-specified topological predicates and the temporal restrictions use Allen’s interval predicates: *before*, *meets*, *overlaps*, *finished*, *during*, *starts*, and *equals* (Allen 1983). Using combinations of these restrictions, the query processor is able to respond to questions such as: “For each month, which changes occurred in the parcel?”, “Which crimes happened on Friday in the south zone of Rio de Janeiro?” and “What was the path followed by this wolf in July of 2007?”

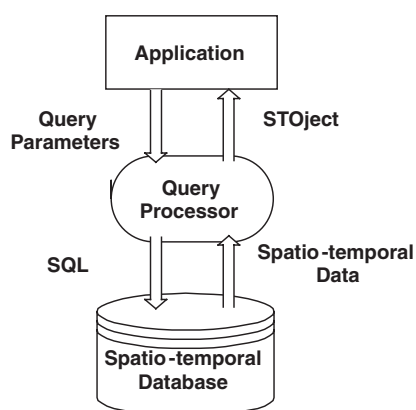


Fig. 12.3 The spatio-temporal data query processor (Source: Ferreira et al. 2005, p. 8)

12.3.6 Spatial Statistics in TerraLib

A GIS produces colour maps of variables such as individual counts quality of life indexes, or company sales in a region. However, to explore the underlying information present in the data, visualisation is not enough. To make effective use of environmental and socioeconomic data, a GIS should provide statistical methods and models. *Spatial statistical* methods measure properties and relations and translate the existing patterns into objective measures. They include geostatistics (Goovaerts 1997), global and local autocorrelation indexes (Anselin 1995), analysis of point patterns (Diggle 2003), regionalization (Openshaw and Albanides 2001; Martin 2003) and spatial regression (Anselin 1988).

One approach to link spatial statistics to GIS is to use loose coupling mechanisms, where the GIS does data conversion and graphic display, and the spatial models run separately. Examples of this approach are the links between SpaceStat and ArcView (Anselin and Bao 1997) and between R and GRASS (Bivand and Neteler 2000). A more recent trend is to integrate spatial statistics methods directly into a GIS. An example of this is GeoDa (Anselin et al. 2006), where a graphical user interface (GUI) is provided for exploratory spatial data analysis on points and polygons.

TerraLib has a basic spatial statistical package, including local and global autocorrelation indexes, non-parametric kernel estimators, and regionalization methods (Assunção et al. 2006). These functions can be used by GIS applications. One such application is TerraView, described in the next section. Additionally, TerraLib provides a direct link with the R programming language using the *aRT* package (Andrade and Ribeiro 2005). R is an open source programming language for statistical computing and graphics and has become a *de facto* standard for developing statistical software (Ihaka and Gentleman 1996).

R has contributors from all over the world, with continuous improvement that incorporates cutting-edge statistical methods. Integration with R can keep a GIS always updated with recent research on spatial statistics. Packages in R relevant to GIS include *geoR* for geostatistics (Diggle and Ribeiro 2007), *splancs* for analysis of point processes (Rowlingson and Diggle 1993), and *sp* that provides general support for spatial analysis (Pebesma and Bivand 2005).

The *aRT* API performs spatial queries and operations in R. It encapsulates TerraLib functions into R objects, and enables R users to read data from a TerraLib database. This coupling satisfies three requirements:

- (a) Statisticians can implement methods of data analysis using R and call TerraLib's facilities for data storage and computational geometry directly for R;
- (b) TerraLib programmers can quickly develop interfaces for calling wrapped R code, which consists of functions and a description of their arguments. They do not need to know about R internals;
- (c) Users of TerraLib-based applications can perform data analysis in R, without knowing R syntax or even without noticing their analysis is executed by R.

An example of the R-TerraLib coupling is shown in Fig. 12.4. The data in this case are a set of point samples stored in a TerraLib database. These data were interpolated into a grid using the *geoR* package (Diggle and Ribeiro 2007) and the result stored as a TerraLib layer and displayed using the TerraView Version 3.0 GIS application.

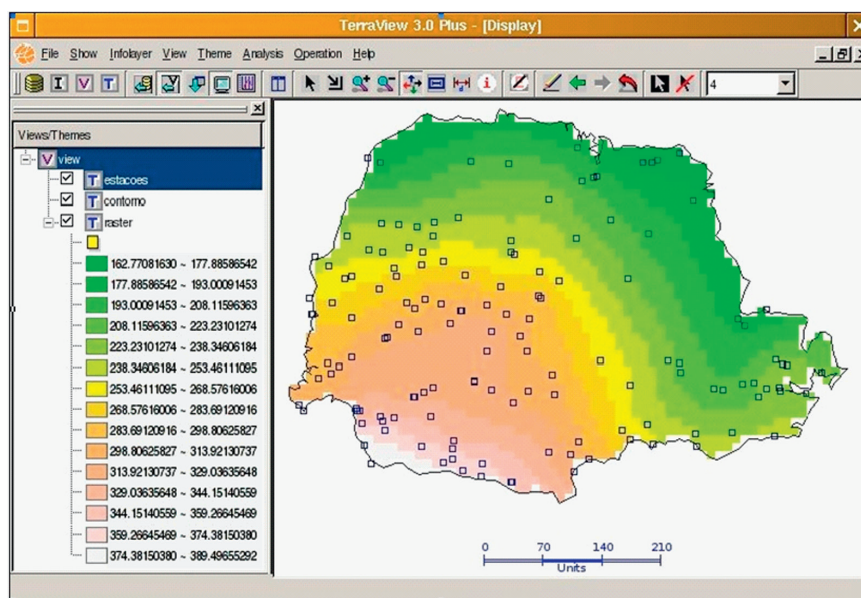


Fig. 12.4 Plotting an R algorithm result in TerraView (Andrade and Ribeiro 2005)

12.3.7 Cell Spaces and Cellular Automata

A cell space is a spatial data type where each cell handles one or more types of attribute. Cell spaces were part of early GIS implementations (Dutton 1978), and later discarded by one-attribute raster data structures, mostly because of efficiency issues. It is time to reconsider this decision and to reintroduce cell spaces as a basic data type in a GIS. Cell-spaces have several advantages over raster-based layers as a means of storing information about continuous spatial entities. Using one-attribute raster data to store results for dynamical models requires storing information in different files. This separation results in increased complexity in data management and user interface design. A cell-space stores all attributes of a cell together, with significant benefits for modelling in contrast to the more cumbersome single value raster approach.

Cell-spaces have been used in the last two decades for simulation of urban and environmental models as part of cellular automata (CA) models (Batty 2000). Most CA models link to a GIS by loose coupling mechanisms, where the GIS performs the data handling and graphic display, and the spatial models run outside the GIS database. This requires extra work for data translation, and may introduce problems of redundancy and consistency. Modeling tools also lack GIS spatial analytical capacities. To address these drawbacks, cell-space models need strong links to the GIS architecture. Using strong coupling, modeling and GIS can be made more robust through their linkage and co-evolution (Parks 1993).

TerraLib supports cell spaces as one of its native data types. It provides functions for storage and retrieval of cell spaces in a DBMS and algorithms for creating cell spaces from vector data. The use of cell spaces enabled development of the TerraME language, which is an add-on to TerraLib that enables simulation in 2D cellular spaces. It supports multi-scale spatial models, where each scale has a different extent and resolution (Carneiro 2006).

Two important innovations in TerraME are its use of anisotropic spaces (Aguiar et al. 2003) and hybrid automata models (Henzinger 1996). Anisotropic spaces arise when modeling natural and human-related phenomena. For example, land settlers in a new area do not occupy all places at the same time. They follow roads and rivers, leading to an anisotropic pattern. However, most spatial statistical and dynamic modeling techniques fail to incorporate spatial anisotropy, leaving out spatial relations that are variable over space. This leads to a serious challenge in producing models that approximate reality, since most real-life spaces are anisotropic.

A *hybrid automaton* is an abstract model for a system whose behaviour has discrete and continuous parts (Henzinger 1996). It extends the idea of finite automata to allow continuous change to take place between transitions. Adopting hybrid automata in spatial dynamic models allows complex models which include critical transitions. Inside each discrete state, the model variables can change. When a critical value occurs, the model moves from the current state to a new one, which is governed by different equations. For example, consider a model for tropical vegetation that has a critical threshold caused by land use change. Under conditions of small land change, the vegetation follows one growth model. When a critical condition is reached, a different growth model must be used. The use of a hybrid automaton allows modeling the tropical vegetation under two different conditions (Carneiro 2006).

Among the typical applications of TerraME are land change and hydrological models. Figure 12.5 shows an application of TerraME for land use change modeling in the Brazilian Amazon. The scenario considers the possible impact of paving a road between the cities of Porto Velho and Manaus. This road crosses areas of currently pristine tropical forest. The model results provide an estimate of how much increase in deforestation could occur in the region from 1997 to 2020 (Aguiar 2006). These models have proven to be useful for supporting public policies that protect the environment and aim at establishing sustainable development practices.

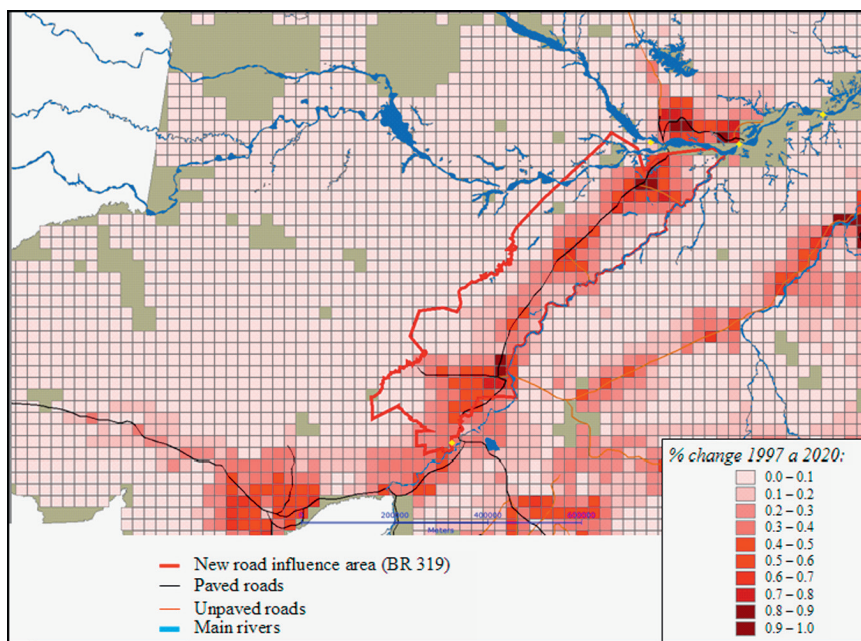


Fig. 12.5 Spatial modeling of projected deforestation of a new road in Amazonia from 1997 to 2020 (Aguiar 2006)

12.4 Development of GIS Applications using TerraLib

This section provides an outline on how to develop a GIS application using TerraLib and describes selected GIS applications that use TerraLib. The general principles of GIS application development are described followed by a description of TerraView, a FOSS4G GIS for spatial data analysis, and TerraAmazon, Brazil's national database for monitoring deforestation.

12.4.1 Building an Application Using TerraLib

TerraLib provides C++ classes that support the higher-level abstractions described in Sect. 12.3, such as *Database*, *Layer*, *View*, and *Theme*. When a programmer builds an application using TerraLib, he/she should use these classes. This section provides a brief guide to the steps involved in GIS application development. TerraLib classes are denoted using monospaced font (e.g., `Database` is the TerraLib class for a spatial database). For more detail, see TerraLib's programming tutorial (Vinhas et al. 2007).

Consider first the case of a GIS application using static data, which has four steps: (a) querying the spatial database; (b) converting the query results (tuples) into objects; (c) manipulating these objects to create new objects; and (d) displaying the resulting objects. To perform these steps in TerraLib, a programmer should include code that does the following:

1. Choose the DBMS that will support the application.
2. Create a TerraLib Database.
3. Connect to the Database.
4. Import data to create a new Layer from standard spatial data formats.
5. Create a view to store a user's view of the database using View.
6. Create a Theme and insert it to the user's View.
7. Define the contents of the Theme, by pointing to the data source (a Layer) and defining attribute and spatial restrictions over that Layer.
8. Load the contents of the Theme.
9. Define the display parameters of the Theme using a Visual.
10. Display the Theme using a GUI toolkit.

A second situation arises when the GIS application uses spatio-temporal data. In this case, following step 8 above, the developer should include code for the following operations:

- 9a. Create a `Querier` (query processor).
- 10a. Define the spatial, attribute and temporal restrictions to apply the query using the `Querier`.
- 11a. Apply the query and get an `STObjectSet` (set of spatio-temporal objects).
- 12a. Manipulate and display the `STObjectSet`.

These steps show that TerraLib abstractions encapsulate a general view of how a GIS works. The abstractions of *database*, *layer*, *theme*, and *view* are especially important, as they provide a link between what is stored in a database and what is selected and manipulated by a user. Thus, a *database* organizes spatial data in *layers* and a user manipulates *themes* according to his/her *view* of the database. Layers store different types of spatio-temporal data, and thus provide the containers needed by the database.

Although these concepts are not part of the current OGC specifications, they arise from decades of experience. Most commercial and FOSS4G applications use them implicitly or directly. Building a user interface on top of these abstractions is simple. A GUI creates events that match to actions that call TerraLib functions. If the user knows the TerraLib ideas, each of these actions will consist of small pieces of code, based on standard examples.

12.4.2 Examples of Open Source Applications

TerraView is a FOSS4G GIS for spatial data analysis, which provides the basic functions of data conversion, display, exploratory spatial data analysis, spatial

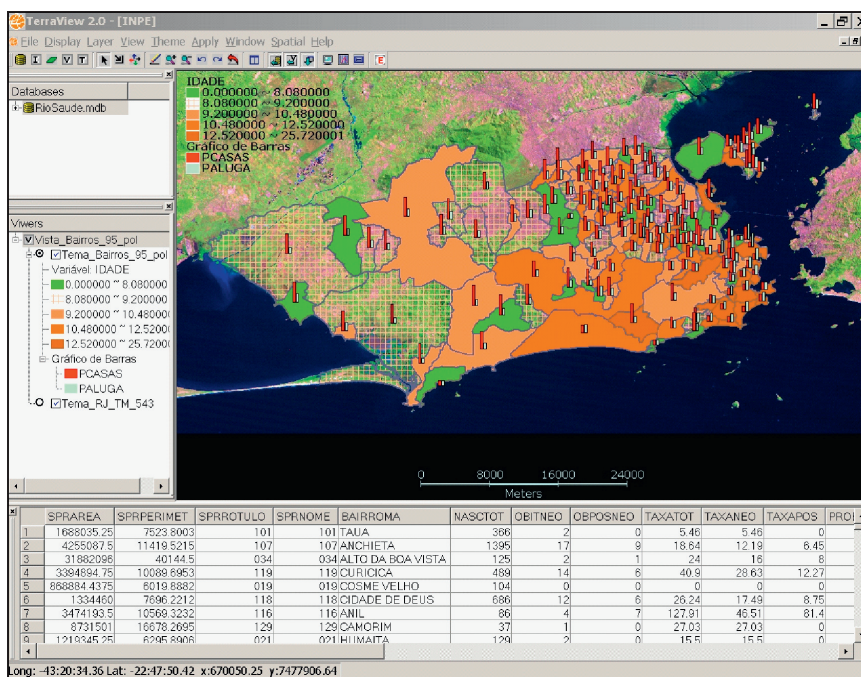


Fig. 12.6 User interface for the TerraView product

statistical modeling, and spatial and non-spatial queries. The project is a general-purpose GIS for TerraLib databases. Many Brazilian public institutions use TerraView for public policy making, including studies in spatial epidemiology and crime analysis. Figure 12.6 shows TerraView's user interface.

TerraView is licensed using the GNU General Public License (GPL), and is available at <http://www.dpi.inpe.br/terraview/>. Its user interface uses the QT cross-platform framework (Blanchette and Summerfield 2006). Programmers can extend its functionalities in two ways. By adapting the menus, they can include new functions or change the behaviour of existing ones. Additionally, TerraView supports plug-ins, which are independent applications that can access a TerraLib database. Using TerraLib, plug-ins have access both to the database and to the display controls (i.e., the lists of views, themes, and the canvas).

A second noteworthy application is *TerraAmazon*, Brazil's national database for monitoring deforestation in Amazonia, developed by INPE and its partner, the Foundation for Space Science, Technology and Applications (FUNCTATE) (<http://www.dpi.inpe.br/terraamazon/>). The DBMS is PostgreSQL version 8.2, hosted on a server running the Linux operating system. The application manages all data workflow by gathering about 600 satellite images and pre-processing, segmenting, and classifying these images for further human interpretation in a concurrent multi-user environment (see Fig. 12.7). The database stores about 2 million complex polygons and grows yearly with 60 gigabytes of full resolution satellite images using

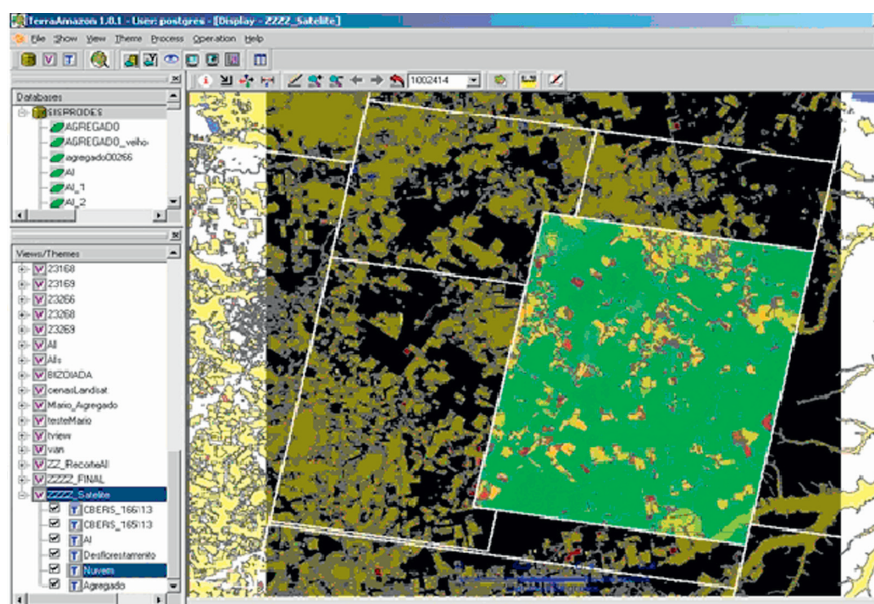


Fig. 12.7 User interface for TerraAmazon

the TerraLib pyramidal resolution schema. A Web site allows seamless display and analysis of full resolution data using TerraLib's PHP extension and TerraLib's OGC WMS server.

12.5 Licensing and Maintenance Policy

One of the important decisions on the TerraLib project was to decide on its license and long-term maintenance policy. The decision considered the nature of the GIS market and the strategy for open source technologies to reach a critical mass of users. The GIS software market is an oligopoly where ESRI, Bentley, and Intergraph have a market share greater than 50% (Daratech 2006). This leads to a "vendor lock-in" effect (Arthur 1994). The "lock-in" effect occurs when a customer is dependent on a vendor for products and services, and cannot move to another vendor without large switching costs. There are many causes for vendor dependence and reluctance to use FOSS4G. First, commercial GIS products use proprietary data formats, making users apprehensive to the costs of data conversion. Second, each GIS adopts a different data model and user interface, which requires training for effective use. Finally, users worry about long-term maintenance of their archives and applications, as the sustainability of FOSS4G projects is often unknown. This results in a conservative policy for most GIS adopters.

Service providers based on OSS face a tough challenge. Convincing a user to change from a commercial to an OS product requires a substantial effort. Users will consider carefully the risks involved in choosing an OS solution, compared

to well-publicized commercial products. Software cost is only part of the problem. Users worry about long-term assurance to protect their investments in data capture and in specialized applications. To convince prospective customers, service providers using a FOSS4G approach need to build custom applications fast and reliably, and this is a task that needs investment. Most service providers consider these applications as IP that needs protection. Thus, they are unwilling to invest in OSS that has binding licenses, such as the GNU GPL.

When deciding on the TerraLib license, the developers considered there should be a strong incentive for commercial companies to use the library to reduce the “vendor lock-in” effects of the GIS market in Brazil. Thus, it was decided to release TerraLib according to the GNU Lesser General Public License (LGPL). The LGPL allows private companies to build proprietary applications on top of TerraLib, and market them as proprietary software, while the TerraLib software itself remains publicly licensed. A second consideration involves development and maintenance of the TerraLib kernel. The Brazilian government has guaranteed its long-term support for the core team of developers. INPE provides capacity building for developers, and supports service companies that use the software.

At the time of writing, there are approximately 10 Brazilian service providers using TerraLib for building commercial applications. Each company has a market focus that includes utilities, the oil industry, agriculture, urban cadastre, and the military. There is evidence that this strategy is paying off. Companies offering GIS services based on TerraLib form 10% of the service provider market in Brazil. This impact on the commercial market is an indicator of a decrease in the “vendor lock-in” effect. The library’s licensing and maintenance policy are arguably an essential part of this result.

An example of a proprietary application that uses TerraLib is InfoPAE, developed by the Computer Graphics Group (TecGraf) of the Catholic University in Rio de Janeiro (PUC-RIO) in partnership with Petrobras (Petróleo Brasileiro S.A.), a Brazilian oil company. The application was designed for emergency response within the oil industry. InfoPAE works with local emergency action plans (LEAPs) that handle significant events. A LEAP is an organized collection of actions, similar to a workflow, coupled with information stored in geographical as well as conventional databases. LEAP frameworks are useful to design large emergency plans. InfoPAE is being used in more than 100 installations of Petrobras in Brazil.

12.6 Conclusion

The design and implementation of TerraLib serve as an example of the challenges involved in building a FOSS4G library that allows innovative applications and supports large-scale applications. One of main lessons learned is that the current set of OGC specifications is not enough to support these goals. Thus, TerraLib has introduced extra abstractions to reduce the cognitive distance between GIS developers and the outcomes of their work. These extra abstractions come at a price. When a FOSS4G developer adheres strictly to the OGC specifications, his/her code runs in

all OGC-compliant libraries. However, adopting TerraLib reduces GIS application development effort, while increasing the cost of using abstractions not supported by other products. OGC-compliant applications will be able to access the part of a TerraLib database that contains OGC's simple features. However, the extra relations used by TerraLib to handle abstractions such as *theme* and *view*, and data structures such as *cell spaces* are invisible to these applications.

Several choices had to be made when introducing solutions for spatio-temporal queries, cell spaces, and raster data handling. Only further experience will show if these were the correct decisions. Another difficult issue concerns the software architecture and design for extensibility. The extensive use of design patterns in TerraLib suits experienced programmers who are comfortable with ideas such as "factory" and "strategy". Novice developers need at least six months training in C++ before they can become skilful in these concepts.

To conclude, developing TerraLib has shown how difficult it is to design a GIS library that combines simplicity and expressiveness. The developers opted for expressiveness at the expense of simplicity in this case. This choice may limit the rate of adoption of TerraLib by the FOSS4G community. Nevertheless, it is the developers' hope the library's assets may be attractive to other developers that want to build large-scale GIS applications.

Acknowledgments *TerraLib's* core team, apart from the authors, includes Laercio Namikawa and Emiliano Castejon at INPE. *TerraView* was designed and implemented by Juan Pinto de Garrido and Lauro Hara. Additional contributors to TerraLib include Tiago Carneiro, designer of *TerraME*, Pedro Andrade, who developed *ART*, Ana Paula Aguiar, who wrote code for cell spaces, and Felipe Castro da Silva and Thales Korting, who developed image processing functions. The technical support of Julio D'Alge for cartographical projections and Leila Fonseca for image processing has been important. We also have important contributions from Paula Frederick, Marcelo Metello, Natacha Barroso, and Leone Pereira Masieiro at PUC-Rio, and Rui Mauricio Gregório and Vanildes Ribeiro at FUNCATE. The TerraLib project is partially financed by CNPq grant no. 552040/02-9. Gilberto Câmara's research is also financed by CNPq grant no. 300557/96-5. The project has also received financial support from FAPESP (Fundação de Amparo à Pesquisa no Estado de São Paulo).

TerraLib code relies on a number of OSS packages. INPE and TerraLib development team thanks the OS community for their efforts. These libraries are: (a) the *zlib* library to compress data when storing raster data in a TerraLib database; (b) the independent JPEG Group's library for JPEG image compression; (c) *libgeotiff* to decode/encode raster data in TIFF/GEOTIFF format; (d) *shapelib* to decode/encode vector data in shapefile format.

As of late-2007, TerraLib consists of 280,000 lines of C++ code and 170,000 lines of third-party open source utilities. The development started in 2001, with an effort of 60 man-years spent so far in TerraLib and TerraView. The library and associated applications may be obtained from the website <http://www.terralib.org>.

References

- Aguiar A, Câmara G, Cartaxo R (2003) Modeling spatial relations by generalized proximity matrices. V Brazilian Symposium in Geoinformatics – GeoInfo 2003, Campos do Jordão, SP, Brazil
- Aguiar APD (2006) Modeling Land Use Change in the Brazilian Amazon: Exploring Intra-Regional Heterogeneity. PhD Thesis, Remote Sensing Program. Sao Jose dos Campos, INPE

- Aksoy S, Koperski K, Tusk C, Marchisio G (2004) Interactive training of advanced classifiers for mining remote sensing image archives. ACM International Conference on Knowledge Discovery and Data Mining, Seattle, WA, ACM
- Alexandrescu A (2001) Modern C++ design: Generic programming and design patterns applied. Addison-Wesley, Reading
- Allen JF (1983) Maintaining knowledge about temporal intervals. *Commun ACM* 26:832–843
- Almeida CM, Batty M, Monteiro AMV, Câmara G, Soares-Filho BS, Cerqueira GC, Pennachin CL (2003) Stochastic cellular automata modeling of urban land use dynamics: empirical development and estimation. *Comput Environ Urban Syst* 27:481–509
- Anselin L (1988) Spatial econometrics: methods and models. Kluwer, Dordrecht
- Anselin L (1989) What's special about spatial data: Alternative perspectives on spatial data analysis. Santa Barbara, CA, NCGIA Report 89-4
- Anselin L (1995) Local indicators of spatial association – LISA. *Geogr Anal* 27:91–115
- Anselin L, Bao S (1997) Exploratory spatial data analysis linking Spacestat and ArcView. In M Fischer and A Getis (eds.) *Recent Developments in Spatial Analysis*, Springer Verlag, Berlin.
- Anselin L (1999) Interactive techniques and exploratory spatial data analysis. In: Longley P, Goodchild M, Maguire D, Rhind D (eds) *Geographical Information Systems: principles, techniques, management and applications*. Geoinformation International, Cambridge
- Anselin L, Syabri I, Kho Y (2006) GeoDa: An introduction to spatial data analysis. *Geogr Anal* 38:5–22
- Andrade PR, Ribeiro PJ (2005) A process and environment for embedding the R Software into TerraLib. VII Brazilian Symposium on Geoinformatics (GeoInfo 2005), Campos do Jordao, Brazil, INPE/SBC
- Arthur B (1994) Increasing returns and path dependence in the economy. The University of Michigan Press, Ann Arbor, MI
- Assunção R, Neves M, Camara G, Freitas CDC (2006) Efficient regionalisation techniques for socio-economic geographical units using minimum spanning trees. *Int J Geogr Inf Sci* 20:797–812
- Austern M (1998) Generic programming and the STL: Using and extending the C++ standard template library. Addison-Wesley, Reading, MA
- Batory D, Singhal V, Sirkin M, Thomas J (1993) Scalable software libraries. *SIGSOFT Softw. Eng. Notes* 18:191–199
- Batty M (2000) GeoComputation using cellular automata. In: Openshaw S, Abraham RJ (eds) *GeoComputation*, Taylor & Francis, London, 95–126
- Blanchette J, Summerfield M (2006) C++ GUI programming with Qt 4. Prentice Hall, Indianapolis, Indiana
- Blaschke T, Hay G (2001) Object-oriented image analysis and scale-space: theory and methods for modeling and evaluating multiscale landscape structure. *Int Arch Photogramm Remote Sens* 34:22–29
- Bivand R, Neteler M (2000) Open source geocomputation: using the R data analysis language integrated with GRASS GIS and PostgreSQL data base systems. 5th International Conference on GeoComputation, Greenwich, UK
- Burrough P (1998) Dynamic modelling and geocomputation. In: Longley P, Brooks S, McDonnell R, Macmillan B (eds) *Geocomputation: A Primer*. John Wiley, New York
- Câmara G, Souza R, Freitas U, Garrido J (1996) SPRING: Integrating remote sensing and GIS with object-oriented data modelling. *Comput Graph* 15:13–22
- Câmara G, Souza RCM, Pedrosa BM, Vinhas L, Monteiro AMV, Paiva JAC, Carvalho MT, Raoult B (2001) Design patterns in GIS development: the TerraLib experience. III Simpósio Brasileiro de Geoinformática, Rio de Janeiro, RJ
- Carneiro T (2006) Nested-CA: a foundation for multiscale modeling of land use and land change. Computer Science Department. Sao Jose dos Campos, INPE. Doctorate Thesis in Computer Science

- Chang SK, Yan CW, Dimitroff D, Arndt T (1988) An intelligent image database system. *IEEE Trans Software Eng* 14:681–688
- Couclelis H (1997) From cellular automata to urban models: New principles for model development and implementation. *Environ Plann B* 24:165–174
- Daratech (2006) GIS markets and opportunities 2006 survey. Cambridge, MA, Daratech Inc
- DeWitt D, Kabra N, Luo J, Patel J, Yu J-B (1994) Client-server paradise. VLDB Conference, Santiago, Chile
- Diggle P (2003) *Statistical analysis of spatial point patterns*. 2nd edn Edward Arnold, London
- Diggle P, Ribeiro PJ (2007) *Model-based geostatistics*. Springer, Heidelberg
- Dutton G (ed) (1978) *First international advanced study symposium on topological data structures for geographic information systems*. Addison-Wesley, Reading, MA
- Elliott J, Eckstein R, Loy M, Wood D, Cole B (2002) *Java swing*. O'Reilly Press, Sebastopol, CA
- Erwig M, Schneider M (2002) Spatio-temporal predicates. *IEEE Trans Knowl Data Eng* 14:881–901
- Feitosa F, Camara G, Monteiro AM, Koschitzki T, Silva MS (2007) Global and local spatial indices of urban segregation. *Int J Geogr Inf Sci* 21:299–323
- Ferreira KR, Queiroz G, Paiva JA, Souza RC, Câmara G (2002) A software architecture for building spatial databases with object-relational DBMS. XVII Brazilian Symposium on Databases, Gramado, RS
- Ferreira KR, Vinhas L, Queiroz GR, Câmara G, Souza RCM (2005) The architecture of a flexible querier for spatio-temporal databases. VII Brazilian Symposium in Geoinformatics, Campos do Jordao, Brazil
- Fonseca F, Egenhofer M, Agouris P, Camara G (2002) Using ontologies for integrated geographic information systems. *Trans GIS* 6:231–257
- Fotheringham AS, Brunsdon C, Charlton M (2002) *Geographically weighted regression: The analysis of spatially varying relationships*. Wiley, Chichester
- Fowler GS, Korn DG, Vo K-P (1995) *Principles for writing reusable libraries*. Proceedings of the 1995 Symposium on Software reusability. Seattle, Washington, United States, ACM Press
- Galton A (2004) Fields and objects in space, time, and space-time. *Spat Cogn Comput* 4:39–68
- Gamma E, Helm R, Johnson R, Vlissides J (1995) *Design patterns: Elements of reusable object-oriented software*. Addison-Wesley, Reading, MA
- Goodchild ME (2003) Geographic information science and systems for environmental management. *Ann Rev Environ Resour* 28:493–519
- Goovaerts P (1997) *Geostatistics for natural resources evaluation*. Oxford Univ Press, New York
- Grenon P, Smith B (2003) SNAP and SPAN: Towards dynamic spatial ontology. *Spat Cogn Comput* 4:69–104
- Gütting RH, Schneider M (2005) *Moving objects databases*. Morgan Kaufmann, New York
- Henzinger TA (1996) The theory of hybrid automata. Proceedings of the 11th Symposium on Logic in Computer Science (LICS'96), IEEE
- Hornsby K, Egenhofer M (2000) Identity-based change: A foundation for spatio-temporal knowledge representation. *Int J Geogr Inf Sci* 14:207–224
- Ihaka R, Gentleman R (1996) R: A language for data analysis and graphics. *J Comput Graph Stat* 5:299–314
- Krasner GE, Pope ST (1988) A cookbook for using the model-view controller user interface paradigm in Smalltalk-80. *J Object-Oriented Program* 1:26–49
- Krueger CW (1992) Software reuse. *ACM Comput Surv* 24:131–183
- Martin D (2003) Extending the automated zoning procedure to reconcile incompatible zoning systems. *Int J Geogr Inf Sci* 17:181–196
- Medak D (2001) Lifestyles. In: Frank AU, Raper J, Cheylan J-P (eds) *Life and Motion of Socio-Economic Units*. ESF Series. Taylor & Francis, London
- Meinel G, Neubert M (2004) A comparison of segmentation programs for high resolution remote sensing data. *Int Arch Photogramm Remote Sens* XXXV:1097–1105
- Meyer B (1990) Lessons from the design of the Eiffel libraries. *Commun ACM* 33:68–88

- Mockus A, Fielding R, Herbsleb J (2002) Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology* 11
- Navulur K (2006) *Multispectral image analysis using the object-oriented paradigm*. CRC Press, Boca Raton, CA
- Openshaw S, Albanides S (2001) Designing zoning systems for representation of socio-economic data. In: Frank A, Raper J, Cheylan J (eds) *Time and Motion of Socio-Economic Units*, Taylor and Francis, London
- Parks BO (1993) The need for integration. In: Goodchild MJ, Parks BO, Steyaert LT (eds) *Environmental modelling with GIS*. OUP, Oxford, 31–34
- Pebesma E, Bivand R (2005) Classes and methods for spatial data in R. *R News* 5:9–13
- Pedrosa B, Câmara G, Fonseca F, Souza RCM (2002) TerraML – A cell-based modeling language for an open-source GIS library. II International Conference on Geographical Information Science (GIScience 2002), Boulder, CO, 2002
- Queiroz GR (2003) *Algoritmos Geométricos para Bancos de Dados Geográficos: Da Teoria à Prática na TerraLib (Geometric Algorithms for Spatial Databases: From Theory to Practice in TerraLib)*. Computer Science. São José dos Campos, INPE. **MSc**
- Rodrigues VL, Andrade MVA, Queiroz GR, Magalhães M (2006) An efficient map overlay algorithm for TerraLib. VIII Brazilian Symposium on GeoInformatics, GeoInfo2006, Campos do Jordão, SP, Brazil, INPE
- Rodrigues VL, Cavalier AP, Andrade MVA, Queiroz GR (2005) Exact algorithms for map manipulation in TerraLib. VII Brazilian Symposium on GeoInformatics, GeoInfo2005, Campos do Jordão, SP, Brazil, INPE
- Rowlingson B, Diggle P (1993) SplanCs: spatial point pattern analysis code in S-Plus. *Comput Geosci* 19:627–655
- Silva MPS, Camara G, Souza RCM, Valeriano D, Escada MIS (2005) Mining patterns of change in remote sensing image databases. The Fifth IEEE International Conference on Data Mining, New Orleans, Louisiana, USA
- Sistla AP, Wolfson O, Chamberlain S, Dao S (1997) Modeling and querying moving objects. *Proceedings of the Thirteenth International Conference on Data Engineering* 422–432
- Veldkamp A, Fresco L (1996) CLUE: A conceptual model to study the conversion of land use and its effects. *Ecol Model* 85:253–270
- Vinhas L, Ferreira KR, Ribeiro G (2007) TerraLib programming tutorial. São José dos Campos, Brasil, INPE (available on <http://www.terralib.org>)
- Vinhas L, Queiroz GR, Ferreira K, Câmara G, Paiva JA (2002) Generic programming applied to GIS algorithms. IV Brazilian Symposium on Geoinformatics, Caxambu, Brazil
- Vinhas L, Souza RCM, Câmara G (2003) Image data handling in spatial databases. V Brazilian Symposium on Geoinformatics, Campos do Jordão, Brazil
- Warmerdam F (2007) Shapefile C Library V1.2, <http://shapelib.maptools.org/>Last accessed July 28th, 2008
- Weber S (2004) *The success of open source*. Harvard University Press, Cambridge, 75
- Worboys M (2005) Event-oriented approaches to geographic phenomena. *Int J of Geogr Inf Syst* 19:1–28