A geotemporal role-based authorisation system

Vijayalakshmi Atluri

MS/IS Department and CIMIC,
Rutgers University,
Newark, NJ 07102, USA
E-mail: atluri@cimic.rutgers.edu

Soon Ae Chun*

Business Department,
City University of New York,
Staten Island, NY 10314, USA
E-mail: chun@mail.csi.cuny.edu
*Corresponding author

Abstract: Geospatial databases include any data with reference to geo-coordinate information. The geospatial data can either be digital raster images that represent the data on the earth in the form of pixels or digital vector data that is primarily from satellites. Due to the fact that many of the high-resolution satellites are commercial in nature, uncontrolled dissemination of the high resolution imagery may cause severe threats to national security as well as personal privacy. The severity of the threats is even more significant when this information is combined with vector maps or other publicly available vector data. In this paper, we present a GeoSpatial Authorisation System (GSAS), which is based on a GeoSpatial Authorisation Model (GSAM), for specifying and enforcing access control policies that makes reference to the spatial regions and locational credentials. The specification of authorisations is based on the spatial and temporal attributes associated with the image data, resolution of the images, geospatial credentials associated with users and privilege modes including view, zoom-in, overlay, view-thumbnail, view-annotation, identify, animate and fly-by that are relevant for geospatial image data. We present the GSAS system and its functionalities.

Keywords: access control; authorisation systems; geospatial databases; information and computer security.

Reference to this paper should be made as follows: Alturi, V. and Chun, S.A. (2007) 'A geotemporal role-based authorisation system', *Int. J. Information and Computer Security*, Vol. 1, No. 1/2, pp.143–168.

Biographical notes: Dr Vijayalakshmi Atluri is a Professor of Computer Information Systems in the MSIS Department, and Research Director for the Centre for Information Management, Integration and Connectivity (CIMIC) at Rutgers University. She received a PhD in Information Technology from George Mason University, USA. Her research interests include information systems security, databases, workflow management, spatial databases, multimedia and distributed systems. She has published over 90 technical papers in the refereed journals and conference proceedings in these areas and is the co-author of the book, *Multilevel Secure Transaction Processing*, Kluwer Academic Publishers (1999). Currently, she serves as a Member of the Steering

Committee for ACM Special Interest Group on Security Audit and Control (SIGSAC) and for the ACM Symposium on Access Control Models and Architectures (SACMAT), General Chair for the 2005 ACM Conference on Computer and Communications Security (CCS), and Co-general chair for the 2005 International Conference on Web Information Systems Engineering. She served as General Chair for 2004 CCS, and Programme Chair for the 2003 CCS, 2000 ACM Workshop on Role-Based Access Control, as Program Co-chair for the 1999 IFIP WG11.3 Working Conference on Database Security, and on the programme committees of a number of conferences and workshops. In 1996, she was a recipient of the National Science Foundation CAREER Award to investigate issues related to incorporating multilevel security into database management systems for advanced application domains such as office information systems, CAD/CAM and workflow systems. In 1999, she received the Rutgers University Research Award for untenured faculty for her outstanding research contributions. Dr. Atluri is a member of the IEEE Computer Society, the Association for Computing Machinery and IFIP WG11.3.

Soon Ae Chun is an Assistant Professor of Information Systems at the City University of New York, College of Staten Island. She received her PhD in Information Technology from Rutgers University. Her research interests include knowledge-based workflow composition and customisation, information systems security and privacy, policy-based web service composition, ontologies and the semantic web and geospatial information and satellite image database systems. She is currently investigating service integrations and security issues as well as mobile and GIS-based environmental decision support tools. She has published in the *Journal of Computer Security*, *IEEE Transactions on Dependable and Secure Computing*, *Journal of Distributed and Parallel Databases*, *IEEE Computational Intelligence Bulletin*, *IEEE Transactions on Geoscience and Remote Sensing* and many conference proceedings. She is a member of the IEEE Computer Society, the Association for Computing Machinery (ACM), the Association of Information Systems (AIS) and the Beta Gamma Sigma National Business Honour Society.

1 Introduction

With the internet and the infrastructure of data clearinghouses, geographic data and services have become more widely and easily available over ubiquitous networks. On the same note, data has become easier to distribute, share, copy and alter. For instance, the National Geospatial Data Clearinghouse by USGS (2004), a component of National Spatial Data Infrastructure (NSDI), provides a gateway to search the geospatial data. The geospatial data sets in these large-scale clearinghouses or organisations include the digital raster images, which store image in the form of pixels, and the digital vector data that store image as points, lines and polygons. While raster images include satellite images, digital orthophoto quads and scanned maps, vector images include the maps of vector type (e.g. shape file), digital line graphs or census TIGER data. Other non-image geospatial data sets are spatially referenced data with locational information, which include census data, voter registration, land ownership data, and land use data. Typically, all these geospatial data objects are multidimensional, comprising attributes such as the x and y spatial coordinates (latitude and longitude), the time of capture, the time of its validity and the resolution.

The source of geospatial image data is primarily from satellites. Creation of high-resolution commercial satellite image data repositories and their use for different applications are gaining importance (see www.spaceimaging.com). There are now more than 15 commercial satellites (e.g. IKONOS, ORBVIEW, EROS and QUICKBIRD) that can provide low-cost high-resolution satellite images (with resolution of 1 m or better), which enable viewing of roads, houses, automobiles and aircrafts, and will make it possible to create highly precise digital maps and Three-Dimensional (3D) fly-through scenes.

While high-resolution low-cost satellite imagery allows the users and organisations to enjoy many benefits, the details in the high-resolution images could reveal vital national resources that could be a target of threats and could encourage industrial espionage, terrorism or more cross-border military attacks. Due to the fact that the high-resolution satellites are commercial in nature, uncontrolled dissemination of the high-resolution imagery may pose severe threats to national security as well as personal privacy. The severity of the threats is even more significant when this information is overlaid with a vector road map, coupled with publicly available data.

Combination of the publicly available personal data pools with high-resolution image data, coupled with the integration and analysis capabilities of modern geographic information systems providing geographic keys such as longitude and latitude, can result in a technological invasion of personal privacy. A person can be identified not only by the name and address but also by *visual exposure*. Therefore, in the near future, it may be technically feasible for anyone, in near real-time, to observe, record and measure the outdoor activities of anyone, at any place in the world, from backyard pools to nuclear plants or to military movements. For instance, Google Earth system[1] combines the data of different resolutions (up to 0.5 to 1 foot resolutions for major cities) to allow the users to fly from space to a specific neighborhood, zoom right in by simply typing an address, search for schools, parks, restaurants and hotels, and tilt and rotate the view to see 3D terrain and buildings. It is a matter of linking an individual building to its building plans and its wiring details. The detailed information like this may give greater advantage to a group with malicious intents.

Policies for prohibiting the release of imagery beyond a certain resolution (such as the guidelines provided by the Department of Commerce), notifying when an image crosses an international boundary or when such a request is made, are beginning to emerge. Currently, commercial entities (e.g. Space Imaging) enforce several security policies while distributing images beyond a certain resolution covering a specific region. Moreover, various governments already voiced against the availability of the high-resolution images of critical national security areas (Haines, 2005). Specifically, these security policies are based on the spatial coordinates, the timestamp and the resolution of images and the credentials of the subjects. Currently, such controlled dissemination is being enforced manually.

In this paper, we present a *GeoSpatial Authorisation System* (GSAS), which is based on the *GeoSpatial Authorisation Model* (GSAM) proposed in Atluri and Chun (2004). GSAS allows the organisations, including large-scale geospatial data clearinghouses as mentioned earlier, to specify and enforce access control policies based on spatial, temporal and resolution attributes associated with the data objects and the credentials associated with users. As a result, the system is capable of specifying security policies with varying degrees of granularities, from coarse to fine-grained. It supports novel privilege modes related to image manipulations.

In addition, because of the fact that GSAS considers the spatial and temporal properties associated with both subjects and objects, it can also be employed in such emerging applications as mobile commerce and ubiquitous computing. Typically, in these environments the spatial location of mobile users is often sensitive or is a determinant of data release. For instance, in an emergency situation, only the doctors in a particular location where the patient had an accident is allowed to view specific sensitive patient records. Or an intelligence agent in the city of Baghdad can view combat resources stationed in Iraq. The authorisation of data depends on the geospatial location information of the subjects as well as the geospatial characteristics of the data. In essence, although in this paper we focus on a specific set of geospatial data, namely satellite imagery, our access control system can be made applicable to various domains where the access control makes reference to the locational information of the data or the subjects.

We limit our focus to access control issues and do not attempt for a complete security solution that requires the user or machine authentication as well as secure communications. These complete solutions may be addressed using X.509 standard public certificates or encryption technologies to properly authenticate and transmit the data securely.

This paper is organised as follows. In Section 2, we summarise the related work in this area. In Section 3, we review GSAM. In Section 4, we present the GSAS and provide details of its architectural components, features, the types of user requests, the steps involved in the authorisation evaluation process, features of the access control module and some screen shots of the system. In Section 5, we present our conclusions and ongoing work.

2 Related work

We review two systems, the Microsoft's Terraserver and the Google Earth, since they provide view and zoom-in that are similar to our GSAS. The Microsoft's TerraServer (Barclay, Gray and Slutz, 1999; Barclay et al., 2002) stores and provides thousands of users with simultaneous access to high-resolution aerial, satellite and topographic data via web browsers over the internet protocol. The data is organised as a multimedia data warehouse where 2-m resolution topographic maps covering all of the USA, the 1-m aerial photos covering 30% of the USA, and 1.5-m resolution aerial photos covering other areas. [2] The USA is divided into ten zones, called scenes. In each scene, the SPIN-2 images are mosaicked together to cover a zone and to provide a seamless pan and zoom between the tiles and resolutions of the same theme within a scene. TerraServer supports a fixed number of resolutions from 1/1,024 m per pixel (scale 0) to 4,096 m (scale 22). The scale is related to resolution in meters per pixel by Scale = log2(resolution) + 10. The highest resolution images currently in the database are 1 m per pixel, which is scale 10. Coarser resolutions are derived by subsampling of higher-resolution images.

TerraServer uses the Gazetteer to find the images by geographic name. It contains the names of about 1.5 million places, with many alternate spellings. A user request for images is made by textual (name lookup), vector map-based navigation (Expedia Map is used for this), famous place list, raster image-based navigation or explicit geo-coordinates. Users can view, zoom-in and out, pan and download the retrieved image. In addition, the image retrieval results also include other informations such as image date,

unit conversion, viewing window size control, relative distance from the nearest city and other images covering the same area.

Similarly, Google Earth (Keyhole Inc., 2001; Google Earth, 2004) provides interactive multi-resolution image mapping service through the Keyhole software. A multi-terrabyte model of the world is constructed from 15-m high-resolution photographic imagery around the globe and 0.5 to 2-ft imageries of major cities in the USA[3], 3D elevation, digital maps and more than 4.6 million US business listings. Users can fly and zoom-in from space-level to street-level images anywhere in the world, tilt and rotate the view or add layers of local hospitals, hotels, subway, map a road trip or measure the distance between two points. Unlike traditional mapping technologies, the Keyhole software creates a dynamic 3D interface for geographic information.

Although both TerraServer and Google Earth provide a web-based large-scale image search, delivery and display system, they do not consider the security and privacy issues associated with the high-resolution images, and therefore do not provide any access control functionality, which is the focus of our system.

The basic authorisation model (Castano et al., 1994) has been extended to support negative authorisation, role-based access control, task-based authorisations and temporal authorisations, dealing with more complex data. More recent extensions include the authorisation models to temporal databases (Atluri and Gal, 2002), data warehouses and derived databases (Rosenthal, Sciore and Doshi, 1999, Atluri and Gal, 2002), semistructured data such as WWW and hypertext systems (Samarati, Bertino and Jajodia, 1996), XML databases (Bertino et al., 2000a; Damiani et al., 2000; Kudo and Hada, 2000), digital libraries (Adam et al., 2002), and to new domains such as workflow management systems (Atluri and Huang, 1996), video databases (Bertino et al., 2000b) and mobile databases (Fu and Xu, 2005). This class of work does not support spatial image data. In this paper, an extension of RBAC, called *Geotemporal Role-Based Access Control* (*Geotemporal RBAC*), is presented to provide access control to the geospatial image data.

In Damiani et al. (2003), an access control to XML-based two-dimensional vector graphics formats such as the the World Wide Web Consortium's Scalable Vector Graphics (SVG) standard has been proposed. This allows fine-grained feature protection for controlled storage and distribution of vector graphic shapes, images and text on the Web. The protected objects can be a conceptual identifier or a conceptual type. The spatial relationship between the subject and object such as *inside, together-with* are used to specify the authorisation rules further. Unlike our approach, this is primarily for the vector objects rather than raster images, and the spatial restrictions in the authorisation are limited to relative positional relationships among subjects and objects, rather than absolute positions in the screen or geospatial region.

The research prototype described in Bertino et al. (2005) is primarily for the vector GIS data that needs to be delivered to the authorised individuals. This approach is based on a geographic area called window and thematic or geometric features. It assumes that the window will have spatial objects in one piece (as a single coverage file). As such, it does not address the following cases: (1) in case when only part of the requested object is available in the coverage file (i.e. the requested window spans several files or coverages), which in fact require clipping; (2) in case when several parts of the request object need to be put together from different files, which require tiling of objects. Unlike the vector data that has a single thematic feature, the satellite images carry information of everything

(various themes and shapes, so to speak) within the covered area. Thus, their model is not completely suitable for geospatial satellite image data.

For pervasive and ubiquitous computing environment with mobile devices, the users access control has been controlled using the context information (Hulsebosch et al., (2005); Zhang and Parasher, 2003), such as the user's physical location (GPS location) as well as other relevant contextual information such as vehicle velocity, device and/or network capacities, temperature and time. The location information is used to determine whether the user has certain application service privileges or not. In Bertino et al. (2005), an extension of the RBAC model called GEO-RBAC is presented to handle the access control to the spatial and location-based information based on the mobile user's physical and logical locations. Our approach accommodates the dynamic access control based on location and time, primarily for the spatially referenced data sets.

In addition to the above, there is a different body of research on protecting images, or part of the image. Some of this research has been focusing on developing the authentication techniques to verify image integrity and authenticity due to the risks from the easy digital image exchange and manipulation, primarily using cryptography-based watermarking techniques (Wong, 1998; Celik et al., 2001; Zhang and Xiong, 2004). In these cryptography approach, the data is freely shared but the mechanisms try to detect the authenticity or integrity of the data, while in access control the data is shared only to the authorised users.

To control access by children images on the internet, Wang, Wiederhold and Firschein (1998) proposed an approach to classify online images as objectionable (such as pornographic or inappropriate contents) or benign, and block the objectionable images. The classification algorithm uses a combination of filters based on colour histogram, icons, texture and wavelet-based shape matching. While this work is focused on content-extraction and access control to the extracted content, our system is based on geospatial coordinates of objects (as well as subjects).

3 Geospatial authorisation model

In this section, we extend the GSAM proposed in Chun and Atluri (2000) and Atluri and Chun (2004) suitable for providing controlled access to geospatial data. While traditional access control models allow the specification of authorisations as a triple, $\langle subject^4, object, privilege \rangle$, GSAM allows the authorisation to be specified in terms of spatial extent (area), time and resolution of geospatial objects, geotemporal roles for subjects and privileges to manipulate spatial objects.

GSAM is the first model to specify and enforce access control based on the resolution, spatial and temporal attributes of the images, and based on spatial and temporal attributes of user credentials. In GSAM, an authorisation *a* is specified as a 4 tuple $\langle ce, ge, pr, \tau \rangle$, where *ce* is a *credential expression* denoting authorised subjects, *ge* is a *geotemporal object expression* to denote a set of authorised objects with a permitted area, *pr* is a set of privilege modes denoting the set of allowed operations, and τ is a temporal extent denoting the time interval during which access is granted. The collection of all authorisations is stored as the *Geospatial Authorisation Base,* (*GSAB*).

3.1 Geotemporal roles

Geotemporal roles for subjects are a set of roles with spatial and temporal credentials indicating that each role is associated with a certain valid region and temporal interval, respectively. In other words, while a user may assume a professor role in a traditional RBAC no matter where the user is, a user may assume different geotemporal roles depending on the location and time where a user is positioned. A user can assume the role of a professor in a classroom during the day time or a property owner in his house. The geotemporal roles are important, especially in the ubiquitous and mobile environment where a user may assume different roles in different environment while moving around and changing the locations and time. Thus, access control policies should be able to refer to the locational and temporal positions of a role. For instance, a user travelling from Newark to New York City may have a student role in Newark so he can check out a library book, but he may assume a visitor role in New York city library where he cannot check out a book. The geotemporal roles thus have geospatial and temporal attributes.

In order to capture the spatial and temporal aspect in the geotemporal role, we introduce the notion of *scene*. A geotemporal role is a role in a scene. Each scene may be considered as a contextual role that can be mapped to a list of actual spatial extents and temporal periods. For instance, a New York City scene is mapped to a specific bounding box (we call it 'rectangle') represented with a geocoordinates, latitude and longitude, and a certain height and width, while a fire scene may be mapped to a set of bounding boxes where the fire incident occurs at different times. This conceptualisation allows to capture a set of similar contexts that have different spatial and temporal extents. This will allow to specify subjects like 'all the policemen who are in the fire scene' or 'all the shoppers in the mall during Christmas season', no matter where the fire or the mall may actually be physically located on the map.

We represent a geotemporal role as a pair $<r; sc>$, where r is a traditional role for subjects as in a RBAC role hierarchy, and sc is a scene that can be associated with a set of geospatial and temporal extents. Each sc is organised in a hierarchy in its domain. For example, an incident domain may have scenes like fire, flood and earthquake, while a shopping domain may have scenes of mall, retail-shop, wholesale area, market and so on. Each sc can be instantiated with scene expression such as scene name, or a specific geotemporal extent such as $<label, lt, lg, h,w, [t_b, t_e]>$ where *label* is a descriptive scene name, such as 'New York City', 'mall' or 'fire', $<lt, lg, h,w>$ denotes latitude, longitude, height and width of a bounding box covering a geographic area of the scene during temporal period between t_b and t_e.

Note that the scenes in the geotemporal roles are associated with the earth's geocoordinates and extents, thus representing the absolute spatial location. The scenes, such as a conference room of a building, have spatial coordinates relative to the building's geographic extent. In this paper, we primarily focus on the geo-scenes that have absolute spatial extents that are mapped to geo-coordinates.

We formalise the subject in the authorisation with a credential expression *ce*. The *ce* is a logical expression to specify an authorised subject with a geotemporal role, including spatial and temporal credentials and other credentials that are associated with a geotemporal role (Atluri and Chun, 2004; Adam et al., 2002). Geotemporal roles are organised in a geotemporal role hierarchy as illustrated in Figure 1.

The credential expression, thus, includes the geotemporal role including a scene, that is spatial and temporal extents to specify the authorised subjects from different time periods and places. This allows the authorisation model to capture an authorised subject that distinguishes a policeman from 9 am to 5 pm from a policeman during the night shift. These two roles may have authorisation on different privileges and different set of objects. Similarly, the spatial credentials allow to distinguish the subjects from different areas (e.g. policeman in the highway area vs. a policeman in charge of residential area).

Figure 1 The police force geotemporal role hierarchy

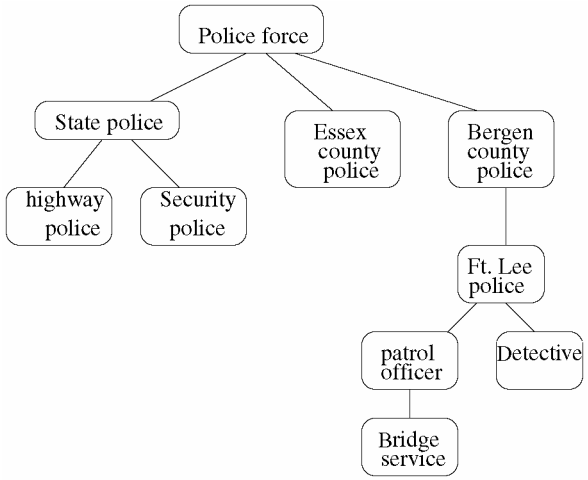

3.2 Geotemporal objects

Similarly, the geotemporal object is defined as an object with a geotemporal scene that maps to an area on the earth's surface. A geotemporal object is specified with a geotemporal object expression *ge*, which is a logical expression of object properties and their values. The object properties are metadata descriptors from object type hierarchy (Atluri and Chun, 2004), such as type of images (IKONOS, Orthophoto, etc.), the geospatial extent of a scene covered by the object (longitude, latitude, width and height), the resolution (the ground area covered by one pixel in the image object) and timestamp (the image download time). Note that the spatial extent in *ge* is a rectangle area expressed either by explicit coordinates (longitude, latitude, width and length), by ZIP codes or highway mile markers or by scene labels such as canonical landmarks such as city or street name.

3.3 Geospatial authorisations

The privilege mode *pr* considers image-specific operations, such as view, zoom-in, overlay, fly-by, and so on. (see Section 4.2.3 for more details.) The τ specifies the valid time period of the authorisation.

Following are some examples of GSAM authorisations:

Given the variable *x* ranging over the subject identifiers and *y* ranging over geospatial object identifiers,

- a_1= ⟨{John(x)}, {type(y)=landsat ∧ rectangle(y)=(50,60,10,10)}, {zoom-in:8}, [1/1/1999, now] ⟩
- a_2= ⟨{NYC-policeman(x) ∧ fire(x)},{type(y)=image ∧ rectangle(y) *overlap* 'New York City' ∧ resolution(y)=1m, type(z)=census-district ∧ data(z)=census92},{view, overlay, identify}, [2/1/2001, 2/1/2005]⟩
- a_3=⟨{Property owner(x) ∧ (home-address(x) *equal* '180 Elm Street, Newark, NJ') ∧ (ownership-period *before* '2000')}, {image(y) ∧ rectangle(y)='180 Elm Street, Newark, NJ' ∧ resolution(y)=1m ∧ timestamp(y) *before* '2000', type(z)=vector ∧ data(z)=property ∧ timestamp(z) *before* '2000', identify} [-∞, 12/31/2004]⟩

Above authorisations can be interpreted as follows: a_1 specifies that John is allowed to access a region centred at point (50,60) with width and height of 10 in landsat images, with a zoom-in level up to 8 during 1 January 1999 and now; The a_2 specifies that a New York City policeman in a fire scene at any location can view the New York City area in 1-m resolution and he is also allowed to overlay census-district map and access census data associated with the map. This authorisation policy is valid from 1 February 2004 up to 1 February 2007. The a_3 specifies that property owners of '180 Elm Street, Newark, NJ' before 2000 are authorised to identify the property information on that address in 1-m resolution images which have been downloaded before 2000. The authorisation is valid until 31 December 2004.

4 The geospatial authorisation system

GSAS is a web-based system[5], whose architecture is shown in Figure 2. As indicated, the administrator can specify subjects, objects and authorisations. We have created a repository of credential types and subject types. GSAS uses Oracle DBMS to store the authorisations, credentials, credential types, object types and image metadata information. The images and the linked tabular data are in the original data format on a Unix file system, and includes formats tiff, gif, jpg, and DBF. Users can submit access requests, which are evaluated by the access control module against the authorisation base, GSAB. The authorised images and data are post-processed (primarily images are chopped, concatenated, composed, animated, etc.) before they are delivered to the user. As such, GSAS provides GUI interfaces for subject, object, authorisation and access request specifications. The web-based user interface, authorisation evaluation and image post-processing are implemented in HTML, JSP, Java Beans and JDBC connection to the database. The prototype runs on Apache Tomcat web server.

4.1 Geospatial image database

Our image database comprises satellite images from NOAA, Landsat images, aerial orthophotos, whose ground resolutions range from 1 km, 28 m, 10 m, 1 m, 8 ft, 4 ft, 2 ft to 1 ft. Much of our image database pertains to the north-east region of the USA, specifically New Jersey, Newark and Hackensack meadowlands area as we capture the

NOAA satellite images daily using the dish antenna located at CIMIC-Rutgers University, which has been serving as a NASA Regional Application Centre (CIMIC Rutgers University, 2002). We have acquired the Landsat images from the NASA archives and the aerial orthophoto images from New Jersey Meadowlands Commission (NJMC)[6].

Figure 2 The system architecture

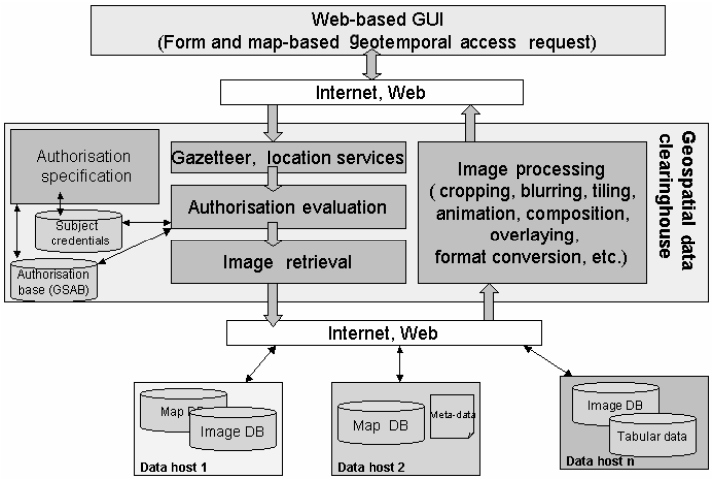

These multi-resolution image data sets are organised in such a way that the lowest resolution images are placed at the top level and the highest resolution at the bottom, forming a so-called *multi-resolution pyramid,* as shown in Figure 3. The region at one level of the pyramid (parent level) is split into four quadrants in the next level (child level). Images in the parent level have lower resolution than those in a child level. Images in the same level have the same resolution. Each quadrant is again split into four subregions in its next level. This process continues at each level of the pyramid, so that the lowest level of the pyramid holds the highest-resolution images. As a result, the entire image database can be visualised as a pyramid. To fit our structure, we have pre-processed the images such that each image belongs to one quadrant only, that is there are no images that cover an area which belongs to two or more quadrants. The images at different levels of the pyramid typically are from different satellite sensors with different resolutions. Note that, our pyramid structure is different from that in the Terraserver system (Barclay et al., 2002). In Terraserver, different resolution levels are generated from a single image, whereas in our pyramid structure each level has images from different satellite sensors with resolutions different from those at other levels.

4.2 Geotemporal role-based specification

The GSAS system comprises interfaces for assigning geotemporal roles to the users and inserting geotemporal role-related credential values into the Credential Base (CB), geotemporal object information into the geospatial object base and authorisation specification.

Figure 3 Multi-resolution pyramid for organising multi-resolution image database

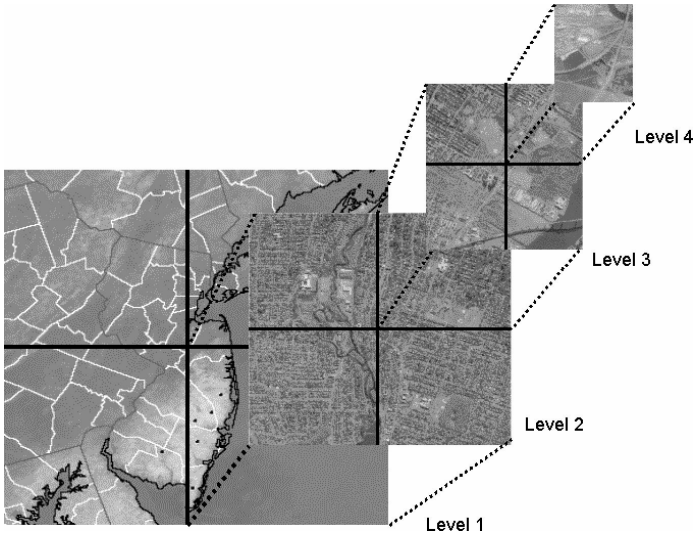

4.2.1 Geotemporal role assignment

The geotemporal roles are organised in a hierarchy as shown in Figure 1. A geotemporal role is associated with a set of credential attributes, called *credentials* (Adam et al., 2002). A geotemporal role inherits the credentials of the geotemporal role of its supertype. The subject specification allows one to assign the users to a geotemporal role and specify values for its associated credentials, and updates the CB. One user (or process), called subject, can be assigned to several geotemporal roles. The set of these credentials and their values for the roles assigned to a subject is denoted as *subject* CB.

Each geotemporal role is associated with a unique identifier and a set of credentials. Each credential is a triple = <*name*, *type, mode*> where *name* is the credential attribute name, *type* is the valid datatype for the credential values and *mode* ∈ {*opt, obl*} denotes whether the credential is optional or obligatory for this credential type. Each subject may hold a set of geotemporal roles, with the scene (i.e. geotemporal) credentials (attribute and value pair) associated with the roles and other credentials associated with the roles. It is important to note that, unlike the traditional credentials, the subject credentials in GSAS consist of both geospatial and temporal (geotemporal) credentials to represent the scene for each role.

Our GSAS subject specification interface allows one to specify a unique user-ID or specific geotemporal roles for each user. As soon as a geotemporal role for a user is assigned, the credentials associated with that role type are automatically displayed so that the subject information can be captured from a drop down list of credential attributes. In our prototype, we have considered geotemporal roles from the domain of property ownership and police authority. These include, person, property owner, policeman, state police, Essex County police and so on. The geotemporal role type hierarchy used in GSAS is shown in Figure 1. Figure 4 shows the interface for subject credential specification. In this figure, since the chosen geotemporal role is 'Property Owner,' shown in Figure 4a, the attributes of this role type, User ID, Name, Age, Address,

Property Address, Lot Numbers and Block Number are displayed as shown in Figure 4b. Selection of another geotemporal role may display different set of attributes. As shown on this screen, the attribute type (numeric or alphanumeric) and whether it is a required attribute or not is also displayed on the screen.

Figure 4 Subject specification (a) Select a geotemporal role, (b) Assign values for credential attributes for a geotemporal role 'Property Owner'

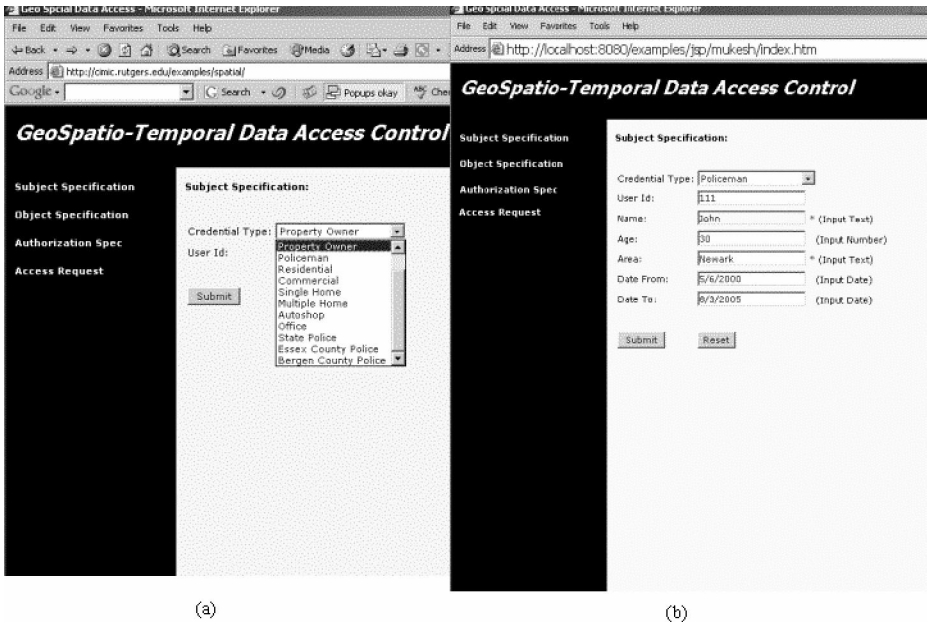

4.2.2 Multi-resolution geotemporal object specification

Objects in GSAS include geospatial raster images with multiple resolutions that represent a geographical region, as summarised in Table 1. In addition, objects include the digital vector data, which store the map and its geographic features as points, lines and polygons, and the tabular data linked to the map, which contains thematic layer information such as census data, voter registration, land ownership data and land use data.

Each geospatial object is associated with a set of metadata, which includes a unique identier, the type of geospatial object, the latitude, longitude, height, width, resolution, timestamp (either image download time or last update time) and the thematic link to the data set associated with the object. Each geospatial object belongs to an object type. We used the following object types in GSAS: raw satellite images (e.g. AVHRR, SPOT, LANDSAT, IKONOS), processed satellite images (e.g. NDVI-NOAA, Composite-Landsat), digital orthophoto quadrangle, aerial photograph, digital elevation model (e.g. DM-raw, DM-hillside, DM-slope), and the type the of features in vector data (e.g. parcels, rivers, parks, schools). Objects types can be organised into a geospatial object type hierarchy. GSAS has implemented the hierarchy shown in Figure 5.

Table 1 Satellite types and their spatial resolutions

Satellites/sensors	*Ground resolution*
NOAA	1.1 km
OrbView-2	1.1 km
MODIS	250–500 m
RadarSat	3–100 m
LandSat	30 m
Aster	15–90 m
Orthophoto	6–7 m
Qickbird	70 cm–2.8 m
IKONOS	1 m
AISA	70 cm–2 m

Figure 5 Geospatial object type hierarchy

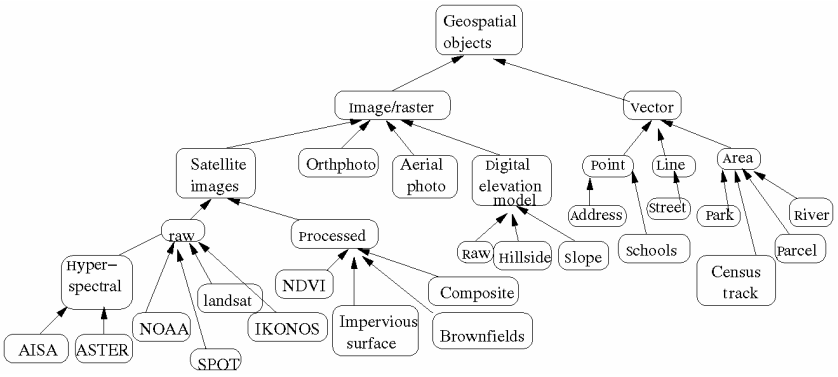

Object specification allows one to insert new objects into the system. As shown in Figure 6, it requires to enter the type of object and, upon selection of an object type from the categories in the object type hierarchy as shown in Figure 6a, the relevant geospatial, temporal, resolution and other attributes are displayed as in Figure 6b. In the figure shown, since NOAA is chosen as an object type, its attributes, Object ID, (MinX, MinY, MaxX, MaxY) that specify the region, Resolution, Time Stamp (object download time), Link and O-Link (the link to the thematic data file associated with the object and the link to the image object file, respectively) are displayed.

Note that our system also allows to specify an area using either explicit geographic coordinates or geographic names (a type of scene). Our geocoding component automatically translates the geographic scenes into corresponding coordinates. Specifically, we have employed a system similar to a geographic gazetteer service[7] to convert place names (i.e. scenes) to coordinates.

We have not built interfaces to create geotemporal role types and object types since they change less frequently. However, our system can easily be extended to accommodate this.

Figure 6 Object specification (a) Select a object type, (b) Assign values for the geotemporal attributes of object type 'satellite'

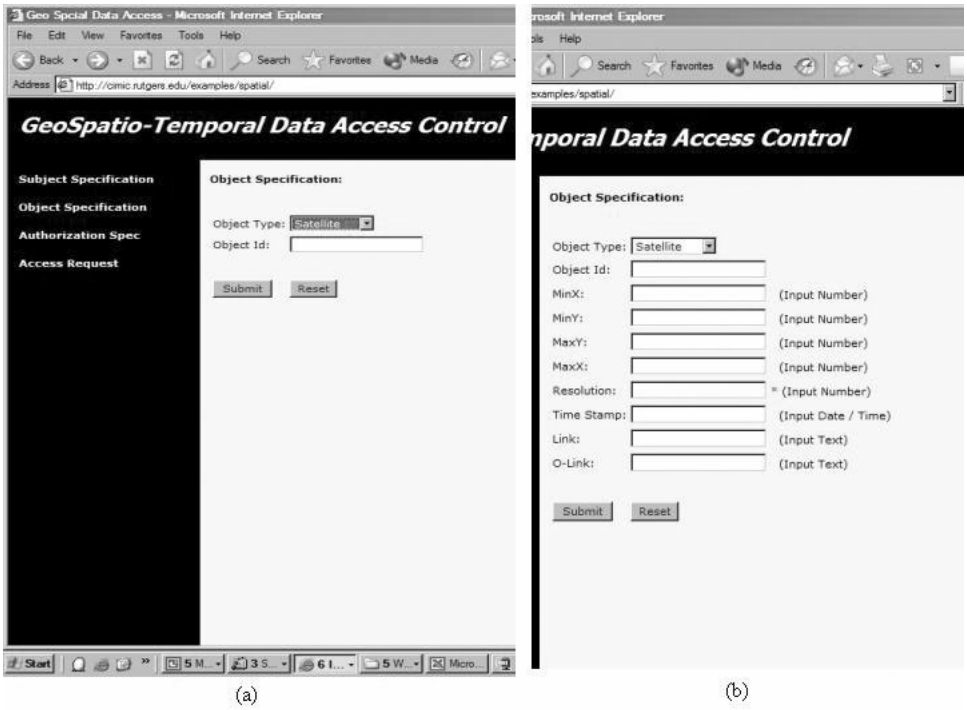

4.2.3 Access authorisations

An authorisation is a 4-tuple <subject, object, privilege, period> that specifies whether a subject (or a subject set) has an access privilege to an object or a set of objects during the period.

Subject: Our system allows the security administrator to specify an authorised subject either by a unique ID or by the geotemporal role and its credential specification, that is a set of attribute and value pairs for a geotemporal role. Once the geotemporal role is selected, its scene and other attributes are automatically displayed to be specified. The geotemporal role type hierarchy information is built and used for the credential reasoning. The obligatory credential attributes are marked such that they are required to be specified with values.

Object: Similarly, the administrator can specify an authorised object either by entering a unique object ID or by choosing an object type with object credentials. The object credentials are specified by entering specific values for the object type-related attributes including spatial and temporal attributes (scene), such as spatial extent, the download timestamp and resolution information. These attributes are displayed automatically using the object type hierarchy.

Privilege: The administrator selects an authorised privilege mode from a drop-down list. GSAS supports a variety of *privilege modes* (permissions). Specifically, in addition

to the conventional operations applied to whole images, the privilege modes include the operations related to image and geographic data manipulations.

The privilege modes are essentially of three types – *viewing, copying* and *maintenance*. The viewing modes include *static* and *dynamic* types. Static viewing modes in turn include view, view-thumbnail and view-annotation, whose purpose is to retrieve data from the data sources and deliver them with basic post-processing operations, such as crop and mosaicking. Dynamic viewing modes include zoom-in, overlay, identify, animate and fly-by, which require, in addition to the basic post-processing operations, geotemporal object integration by building specific modules (in case of animate and fly-by).

View allows a user to see an image object covering a certain geographic area; zoom-in allows a user to view an image covering a certain geographic area at a specific higher resolution; overlay allows the users to generate composite images, where a composite image is constructed from multiple images by first georegistering and then overlaying them one on top of another; identify allows the user to view the tabular data linked to an image; animate allows a user to obtain a time series of images and integrate them to show the changes in the images; and fly-by allows a user to traverse from one location to another a multi-resolution browsing from low-resolution images to high-resolution images or vice versa.

The copying modes, download and download-data, allow the source files to be downloaded. Unlike the text data where the display privilege implies the copying privilege, the viewing and copying are distinguished as separate privileges with geospatial data since the objects displayed on the web browser often are image gif files, but not the original source files. The maintenance modes include insert, delete, update and compose. The users with compose privilege can create and insert value-added images, using images in the database.

Period: The time period when the authorisation is valid is specified with the temporal duration. If it is unspecified, the authorisation is valid for indefinite time period.

Figure 7 shows the interface for authorisation specification. The objects and subjects specified appear as drop-down lists to facilitate easier authorisation specification. Note again that as soon as one selects the credential and object types, the attributes associated with them are displayed on the screen. The different privilege modes are also shown as a drop-down list. Additionally, a time interval can be specified during which the specified authorisation is valid. Instead of specifying the objects with their coordinates, one may choose the 'select city' option (on the right-hand side in Figure 7), where the region covered by the selected city is automatically added to the authorisation.

4.3 Access request specification

In GSAS, an access request is supported primarily in two ways:

1 *Explicit access request*: The user explicitly enumerating the geotemporal characteristics of the requested object, such as location name or geocoordinates, timestamps or resolutions, with explicit credentials.

2 *Implicit access request*: Once the subject gets authenticated and the credentials verified, the request can be made through the click on the location on the map or on a base image, using a map or image interface. There is no need to explicitly specify the image characteristics. As long as the selected location is within the authorised region

and within the resolution level in the multi-resolution pyramid discussed in Section 4.1, the geotemporal objects will be granted to access, level by level, from the low-resolution image to higher-resolution image up to the granted resolution level.

In the following sections, we discuss these two types of access requests and how the system supports both.

Figure 7 Authorisation specification

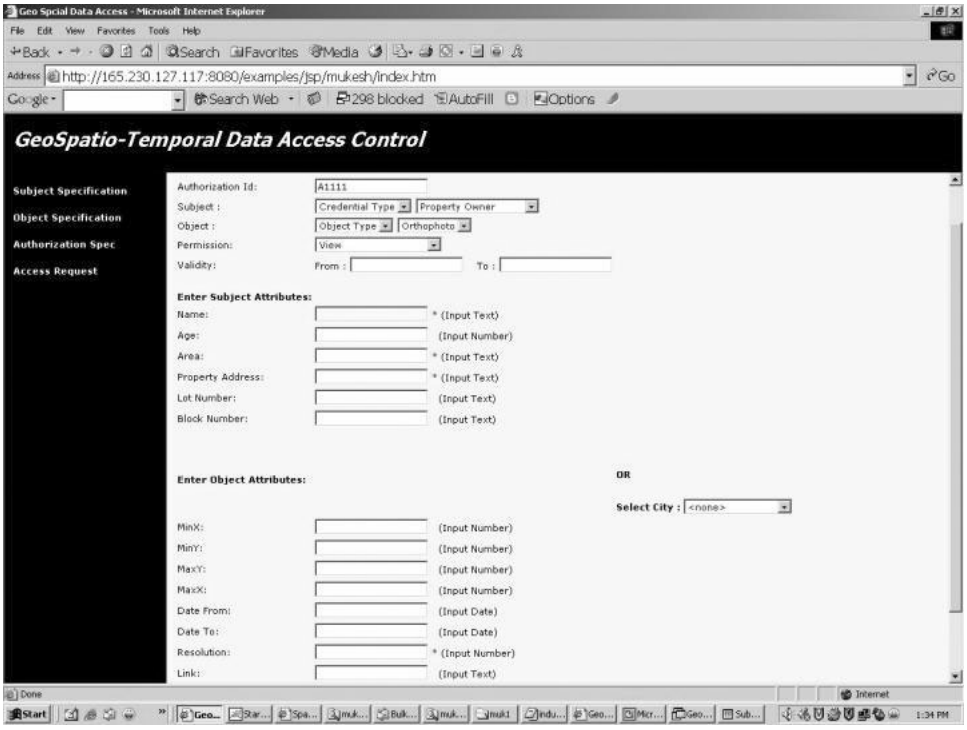

4.3.1 Explicit access request

Explicit user access request, *ur* = <s, re, m>, can be specified in a number of ways. The users can submit access requests using object attributes such as geospatial, temporal and resolution of images, and using subject credential expressions, as shown in Figure 8. The user can also explicitly specify the object identifiers, instead of object characteristics. The object request can also use a scene, such as a canonical geographic name, to specify the region, instead of exact geocoordinates (rectangles), 'Newark, New Jersey'. Our system supports a gazetteer service to geo-code the scene or location name into its corresponding region, as shown in the bottom of Figure 8. Similarly, a user access request can use the user identifier or geotemporal credentials for specifying the geotemporal role.

Following are the examples of different types of user requests.

$ur_1 = \langle$John, {12, 24, 100}(x), view $\rangle$

ur_2 = ⟨Mary, Satellite-Image(x) ^ rectangle=(50, 60, 10, 10) ^ resolution =1m ^ timestamp=[8/1/2001-now] ^ link=property), identify ⟩

ur_3 = ⟨ compute-ndvi, Landsat(x) ^ rectangle= (50, 60, 10, 10) ^ timestamp =[5/1/2000-now], overlay⟩

ur_4 = ⟨ policeman(x) ^ area=(50, 60, 10, 10), SPOT(y) ^ address='Newark, NJ', view ⟩

Figure 8 Explicit access request

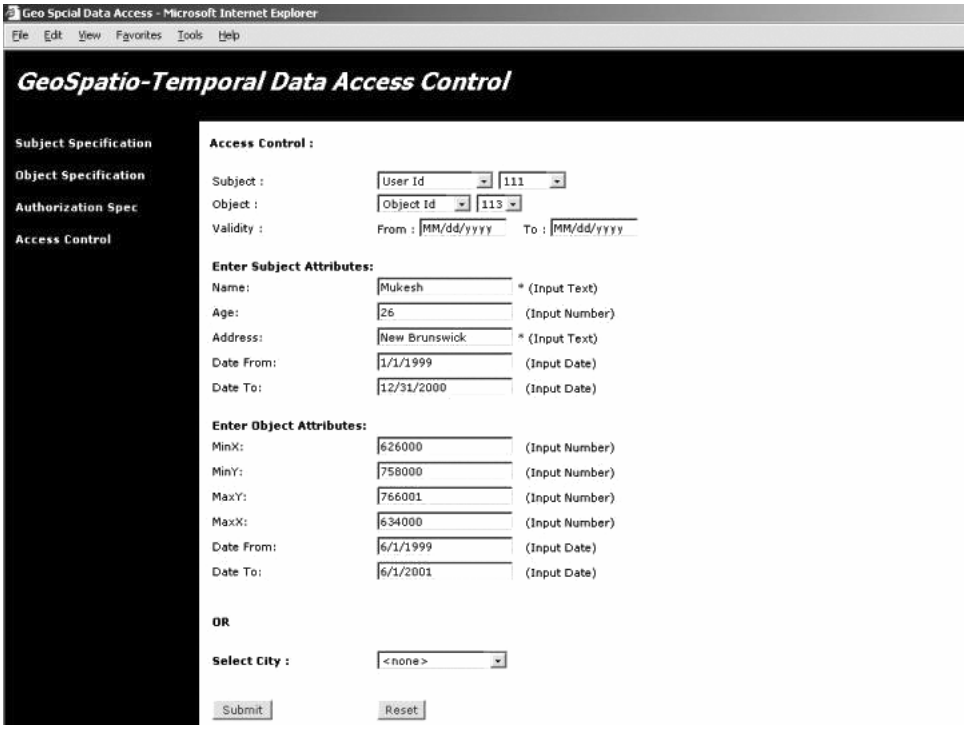

User request ur_1 states that John wants to view spatial data with identifiers equal to 12, 24 and 100. In ur_2, 'Mary requests to retrieve the linked property information of a specific rectangular region represented by (50, 60, 10, 10) from the images of 1-m resolution downloaded between 1 August 2001 and now'. In ur_3, 'an application program *compute-ndvi* is requesting to overlay the Landsat images of a specific region represented (50, 60, 10, 10) downloaded between 1 May 2000 and now'. In ur_4, 'The policeman in the region (50, 60, 10, 10) wants to view SPOT satellite images of Newark, New Jersey area'. Essentially, the requested geospatial objects can be specified in several ways:

1 object identifiers (ID) (as in ur_1), which means the user is requesting the whole area covered by the image identified with IDs

2 a geographic coordinate representing a rectangle (as in ur_2 and ur_3), which specifies a region or an area that may be contained in an image or that may span across multiple images

3 a scene such as a conventional geographic name (as in ur_4), which designates a region or an area. Different cases are shown in Table 2. Similarly, the subjects can be specified in four different cases as shown in Table 3.

Table 2 Different cases of object specification of access request and authorisation

	Requested object (*re*)	*Authorisation object* (*ge*)
Case 1	ID	ID
Case 2	ID	Region
Case 3	Region	ID
Case 4	Region	Region

Table 3 Different cases of subject specification for access request and authorisation

	Requested subject (*s*)	*Authorisation subject* (*ce*)
Case 1	ID	ID
Case 2	ID	Region
Case 3	Region	ID
Case 4	Region	Region

The authorisation verification considers four different cases as shown in Table 2, depending on the spatial reference of the requested and authorised objects.

Case 1: If the requested object and authorisation object are specified with image IDs, then the access control verification checks of the requested images are in the authorised set of image id's.

Case 2: If the spatial reference in the requested object is specified with image ID and the authorised object is specified as a region, then the verification process verifies the entire authorised region that contains the spatial extent of the image ID, identifies the images that contain the authorised region, and matches the ID's with the requested image ID.

Case 3: If the requested object is a region, and the authorised object is specified as image IDs, the verification process checks if the authorised image IDs, spatial extent contains or overlaps with the requested region. If the authorisation objects contain the requested object, then access is granted.

Case 4: If the requested object and authorisation object are both expressed as regions, then the intersecting area of these two regions is computed and the images that cover the intersected area are returned.

4.3.2 Implicit access request

In case of implicit specification, an interactive selection of an area of an image can be done by simply clicking the mouse on a map or an image. In GSAS, we have built a component that translates the area designated by the mouse click into geographic coordinates. Specifically, these include view and zoom-in operations. As soon as a user clicks on a point, the quadrant region containing the point is considered. The system checks if the zoom-in request is allowed for that subject in that selected region (essentially, verifying if the subject is allowed to view an image at that level of resolution). If so, the next higher-level resolution image covering the selected area is presented. Figure 9 shows the GUI interface for this navigational access control, where

the 8-ft resolution true colour aerial image is first zoomed-in to show the image of 4-ft resolution image that is zoomed in to show the 2-ft resolution image and then 1-ft resolution image.

4.4 Processing of access requests

The access control mechanism evaluates a user request *ur* against the authorisation base *GSAB* to determine whether the user request should be granted or denied. Figure 10 shows the evaluation steps. The user request is made using a subject identifier or a credential, object ID or spatial expression and privilege mode. The system evaluates the requested object with the authorisation base GSAB. The requesting subject is verified using the credential base.

Figure 9 GUI for implicit access request for successive zoom-in

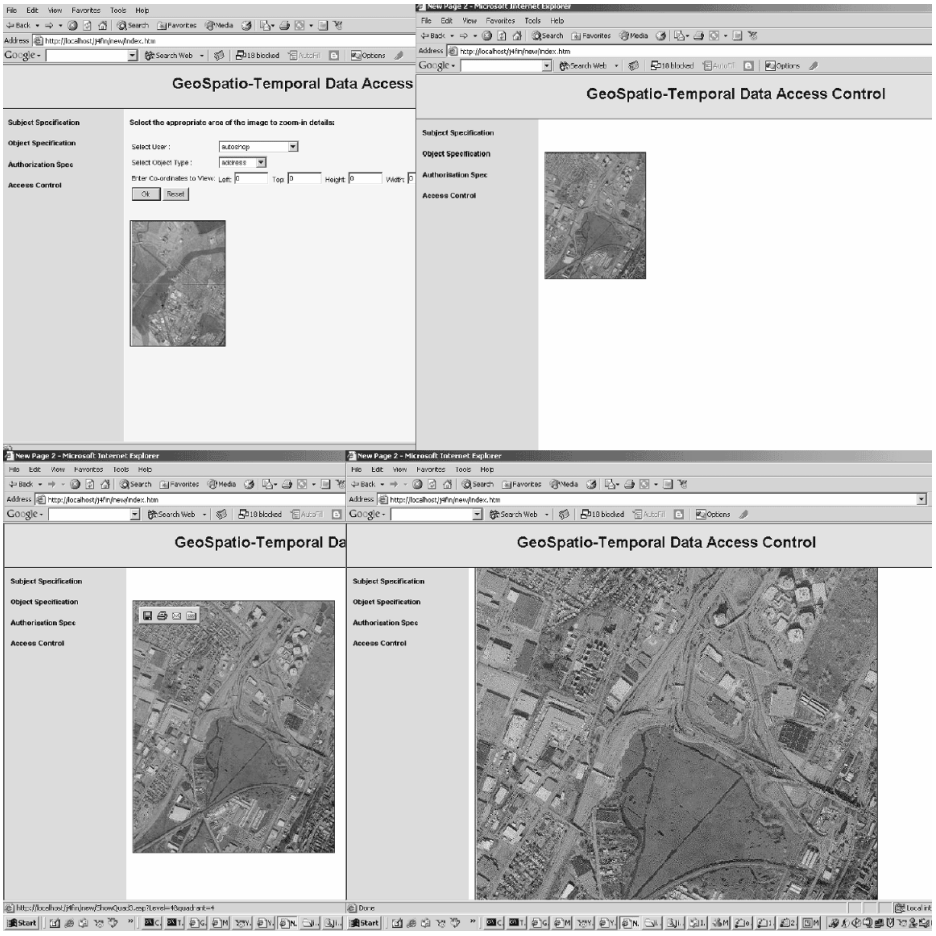

Figure 10 The authorisation verification process

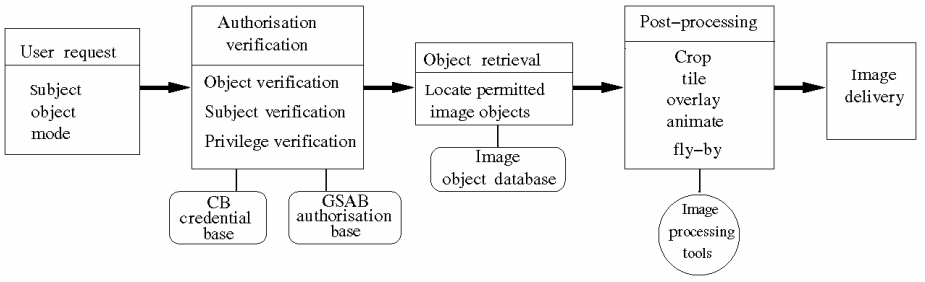

The request mode is verified with the authorised privileges. Once the verification is finished, the authorised image(s) are retrieved from the object repository and fed into the post-processing module, where the images are cropped, tiled, an animated image is created with the images along the time sequence, or a fly-by is created with a starting and ending points, moving from different locations and resolutions of images. The post-processing module uses image processing tools. The image product from the post-processing module is delivered to the user. The following sections explain how each component functions.

4.4.1 Authorisation evaluation

Recall that each Authorisation a is specified as $a = \langle ce, ge, pr, \tau \rangle$. Therefore, the access control mechanism requires to identify the authorisations relevant to *ur* from *GSAB*. To accomplish this, it identifies each authorisation *a* by verifying

1 if the authorised objects specified by *ge* satisfy the requested object characteristics, such as the geographic area, resolution level and temporal extent specified in the user request

2 if the subject specification in *a* match with the credentials of user requesting access and

3 if the privilege mode specified in the access request is found in some authorisation *a* that satisfy conditions (1) and (2) above.

Specifically, the above authorisation verification process can be outlined as consisting of the following steps: (i) *Object Identification*, (ii) *Subject credential evaluation*, and (iii) *Privilege evaluation*. In the following, we describe each step of the evaluation process in detail.

(a) *Object identification*: This identifies the authorisations whose object characteristics (e.g. region, resolution, timestamp) satisfy the object characteristics of the requested object. It needs to consider the different types of user requests, that is requests with image IDs, canonical geographic names and a region. It identifies the authorisation entries whose authorisation area overlaps with the requested area and that meet the temporal, resolution and other requested object characteristics.

1 If the user requests the images with specific IDs, then the evaluation process first finds the spatial extents, rectangle(ID), for each requested image ID. Each authorisation *a* is evaluated against either *id*, if *a* is also specified with the IDs, or *a* is evaluated with the rectangle(*id*) in case *a* is expressed with rectangle(*ge*) in the geospatial expression *ge*.

2 If the user requests the images with region (specified as a rectangle), then the evaluation process matches the spatial extent specified in *ge* with that of the requested region (rectangle). If the authorisation object is expressed in terms of image IDs, then the rectangle of the images specified by this *id* is retrieved and matched with the rectangle specified in the access request.

3 If the user requests the image with geographic name, then the gazetteer service is invoked to convert the geographic name into a region. Then the evaluation process in step (b) is used.

(b) *Subject credential evaluation*: The subject verification requires the activation of geotemporal roles. The user's current location and time of the request may be captured through the location-service provider (using GPS system) and be sent over to the access control evaluation. The scenes for the geotemporal roles assigned for the user in the authorisation base are matched with the user's actual position. The geotemporal roles, in the authorisation base, whose scenes match with the geospatial characteristics of the user's current spatial and temporal position will be activated. In addition, the geotemporal roles of the past scenes (e.g. Newark policemen in 1971) will be also activated for consideration. The geotemporal roles with past scenes need credential verification whether the user was actually a policeman in Newark during 1971[8].

The subject credential evaluation checks whether the user's active geotemporal roles satisfy the credentials for the matching geotemporal roles specified in the authorisation. In other words, it evaluates the user's credentials stored in subject CB against the credential expression specified in the authorisations, yielding the authorisations that satisfy not only the requested objects but also the user's credentials. It checks the temporal and spatial overlap areas between activated roles in the subject specification and authorisations. Subject credentials associated with the user's active geotemporal role IDs are retrieved from the subject *CB*. The system evaluates the authorisations identified in step 1 above with subject credentials.

(c) *Privilege evaluation*: The requested privilege mode is evaluated with the privilege mode specified in authorisations. The authorisations that are satisfied with the requested object and subject from steps 1 and 2 are further verified with the requested mode. This evaluates the request mode with the privilege mode specified in the authorisation specification, and evaluates the various spatial and temporal operators with specific values in the request and the authorisation. Only objects that meet the object, subject and privilege specifications in the authorisation base are collected with the authorised area. It returns the authorised object identifiers along with authorised area for each object that satisfies the subject, object and privilege of the request.

4.4.2 Post-processing and delivery

Once the authorisation is evaluated and the authorised objects are identified along with the authorised area, images are retrieved from the database and sent to the post-processing component. This post-processing and delivery module performs image processing, which include crop, tile, overlay, animate or create a fly-by video etc., according to the requested privilege mode. In the following, we elaborate these functionalities.

(a) *View and zoom-in*: If the requested information is the annotation (metadata) information, the metadata of each authorised objects are retrieved. If the request is to view an image, then the actual image file is retrieved but post-processed to deliver only the authorised area in the image. For instance, when the images are larger than the requested and authorised areas, this post-processing component *crops* so that only the authorised and requested areas are delivered to the user. Figure 11 shows the original image in the right-hand side and the cropped image that contains only the allowed area.

Alternatively, if the authorised area spans several images, the authorised area of each image is *tiled* together as a mosaic and then delivered to the user. If the zoom-in mode is requested, the zoom-in level is considered with the authorised resolution level. In summary, it performs either cropping or tiling, based on the spatial relationships between the requested region and the authorised region. Figure 12 shows the area requested is from four different images tiled together.

(b) *Overlay*: In the case of overlay, the image processing component shows multiple images overlaid one on top of the other, by carefully synchronising the geospatial coordinates of the images. An example may include a high-resolution raster image of the Newark area overlaid with a vector image showing the map of the region with marked roads.

Figure 11 The area allowed for access is cropped from the original image

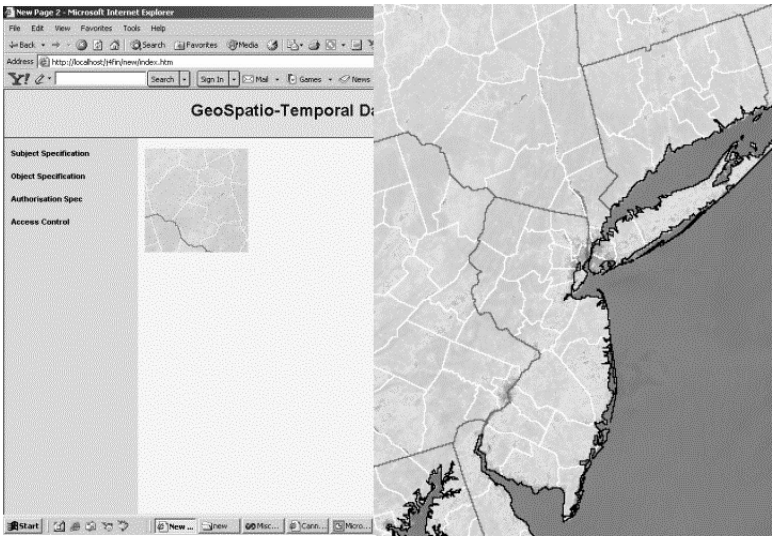

Figure 12 The area allowed for access is tiled from four different images

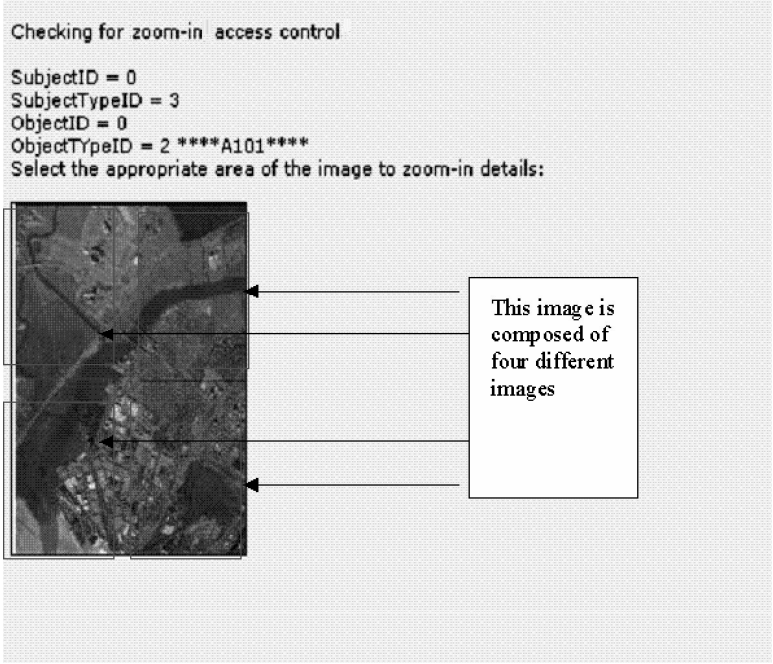

(c) *Animate*: If the access request consists of the animate privilege, then a series of authorised images should be composed together and shown as an animation. A sequential display of the images during a certain time interval shows how the images vary over a period of time. An example may include the variation of vegetation index over 1 year period in the north-east region of the USA. The users are able to specify the rate at which the animation should be displayed.

(d) *Fly-by*: it simulates the fly-by with multiple images in different resolutions by constructing a video of the region where a user specifies the origin and destination by specifying their coordinates and resolutions, as well as the speed of traversal. For example, the origin could be a low-resolution point in Newark and the destination could be a high-resolution point in the Meadowlands district.

5 Conclusions

In this paper, we have presented a GSAS, which is based on the GSAM, for specifying and enforcing the fine-grained access control policies that allow the specification of authorisations based on the spatial and temporal attributes associated with the image data, resolution of the images and credentials associated with the users.

We are currently extending this work along several directions. The first direction deals with improving the response time for processing of access requests. To address this issue, we have been developing unified index structures that are capable of indexing both geotemporal objects and the authorisations that govern access to them. Current GSAM system does not use any spatio-temporal index on authorisations or objects except that

provided by Oracle on the image and authorisation identifiers. We plan to enhance GSAS by incorporating a suitable index and study its performance.

The second extension is to incorporate more flexible user request with an arbitrary area specification with graphical interface, such as drawing a rectangular area or polygon area to zoom-in. The third direction for extension is to allow more content-based retrieval and access control of the satellite images. In order for this, we are also looking into semantic object extraction from the satellite images. Currently, we have post-processed the arbitrary set of images into a fly-by, but semantic combination should be considered to deliver more meaningful products.

Acknowledgements

The work of Atluri is supported in part by the National Science Foundation under grant IIS-0242415 and the Meadowlands Environmental Research Institute (MERI). The work of Chun is partially supported by MERI. We acknowledge Artigas, MERI, New Jersey Meadowlands Commission, for the information on geospatial satellite images, their analysis and processing; Elefante for providing high-resolution images over the New Jersey Meadowlands area; Baumann, Rasdaman.com, for helping with the OpenGIS standards and for providing Rasdaman database; Guo for preprocessing satellite images; Sethi and Shin for implementing GSAS, and Yu and Bansal for building the animate and fly-by components of GSAS.

References

Adam, N.R., Atluri, V., Bertino, E. and Ferrari, E. (2002) 'A content-based authorization model for digital libraries', *IEEE Transactions on Knowledge and Data Engineering,* Vol. 14, pp.296–315.

Atluri, V. and Chun, S. (2004) 'An authorisation model for geospatial data', *IEEE Transactions on Dependable and Secure Computing*, Vol. 1, pp.238–254.

Atluri, V. and Gal, A. (2002) 'An authorisation model for temporal and derived data: securing information portals', *ACM Transactions on Information Systems Security*, Vol. 5, pp.62–94.

Atluri, V. and Huang, W-K. (1996) 'An authorization model for workflows', Paper presented at the *Fifth European Symposium on Research in Computer Security*, Rome, Italy, and Lecture Notes in Computer Science, Springer-Verlag, pp.44–64. In proceedings.

Barclay, T., Gray, J. and Slutz, D. (1999) '*Microsoft terraserver: a spatial data warehouse', Technical report, Microsoft Research*, Advanced Technology Division Microsoft Corporation, Redmond, WA.

Barclay, T., Gray, J., Strand, E., Ekblad, S. and Richter, J. (2002) 'Terraservice.net: an introduction to web services', Technical report, Microsoft Research, Advanced Technology Division, Technical Report MS-TR-2002-53.

Bertino E., Damiani M.L. and Momini D. (2004) 'An access control system for a web map management service', Paper presented at the *14th International Workshop on Research Issues in Data Engineering, Web Services for E-Commerce and E-Government Applications, 28–29 March 2004, Boston, MA*, (pp.33–39). IEEE Computer Society. In proceedings.

Bertino, E., Castano, S., Ferrari, E. and Mesiti, M. (2000) 'Specifying and enforcing access control policies for XML document sources', *World Wide Web Journal*, Vol. 3, pp.139–151.

Bertino, E., Catania, B., Damiani, M.L., and Perlasca, P. (2005) 'Geo-rbac: a spatially aware rbac', Paper presented at the *SACMAT '05*: *the 10th ACM symposium on Access control models and technologies*, pp.29–37, In proceedings.

Bertino, E., Hammad, M.A., Aref, W.G. and Elmagarmid, A.K. (2000b) 'An access control model for video database systems', Paper presented at the *CIKM.* In proceedings.

Castano, S., Fugini, M., Martella, G. and Samarati, P. (1994) *Database Security*. Reading, MA: Addison-Wesley.

Celik, M., Sharma, G., Saber, E. and Tekalp, A. (2001) 'A hierarchical image authentication watermark with improved localization and security', Paper presented at the *IEEE International Conference on Image Processing (ICIP 2001)*, pp.502–505, Thessaloniki, Greece. In proceedings.

Chun, S. and Atluri, V. (2000) 'Protecting privacy from continuous high-resolution satellite surveillance', Paper presented at the *14th IFIP WG 11.3 Workshop on Database Security*, pp.399–420. In proceedings.

CIMIC Rutgers University (2002) NASA Regional Application Centre. Available at: http://cimic.rutgers.edu/rac/.

Damiani, E., di Vimercati, S.D.C., Fernandez-Medina, E. and Samarati, P. (2003) 'An access control system for SVG documents', Paper presented at the *DBSec, IFIP Conference*, Cambridge, UK. In proceedings.

Damiani, E., di Vimercati, S.D.C., Paraboschi, S. and Samarati, P. (2000) 'Design and implementation of an access control processor for XML documents', Paper presented at the *Ninth International World Wide Web Conference (WWW9)*. In proceedings.

Fu, S. and Xu, C-Z. (2005). 'A coordinated spatio-temporal access control model for mobile computing in coalition environments', Paper presented at the *19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05) – Workshop 17*. In proceedings.

Google Earth (2004) Google acquires Keyhole Corp, Google Press Center, Available at: http://www.google.com/press/pressrel/keyhole.html

Haines, L. (2005) South Korea throws strop at Google Earth: military installations laid bare. The Register, Available at: http://www.theregister.co.uk/2005/08/31/google_earth_korea/

Hulsebosch, R., Salden, A., Bargh, M., Ebben, P. and Reitsma, J. (2005) 'Context sensitive access control', paper presented at *SACMAT '05: the tenth ACM symposium on Access control models and technologies*, pp.111–119. In proceedings.

Keyhole Inc. (2001) Keyhole's EarthViewer Technology, System demo presented at Vortex 2001, http://www.keyhole.com//body.php?h=news&t=20010522

Keyhole Inc. (2001) 'Keyhole's EarthViewer technology', System demo presented at Vortex 2001, Available at: http://ww.keyhole.com//body.php?h=news&t=20010522

Kudo, M. and Hada, S. (2000) 'XML document security based on provisional authorization', Paper presented at the *ACM Conference on Computer and Communication Security (CCS 2000)*. In proceedings.

Rosenthal, A., Sciore, E. and Doshi, V. (1999) 'Security administration for federations, warehouses, and other derived data', Paper presented at *IFIP WG 11.3 Thirteenth International Conference on Database Security*, pp.209–223. In proceedings.

Samarati, P., Bertino, E. and Jajodia, S. (1996) 'An authorisation model for a distributed hypertext system', *IEEE Transactions on Knowledge and Data Engineering,* Vol. 8, pp.555–562.

USGS (2004) 'USGS Node of the National Geospatial Data Clearinghouse, Available at: http://nsdi.usgs.gov/,

Wang, J., Li, J., Wiederhold, G. and Firschein, O. (1998) 'System for classifying objectionable websites', Paper presented at the *International Workshop on Interactive Distributed Multimedia Systems and Telcommunication Services,* Vol. 1483, pp.113–124, Springer Verlag. In proceedings.

Wong, P. (1998) 'A public key watermark for image verification and authentication', Paper presented at the *IEEE International Conference on Image Processing,* pp.425–429, Chicago, USA. In proceedings.

Zhang, G. and Parasher, M. (2003) 'Dynamic context-aware access control for grid applications', Paper presented at the *4th International Workshop on Grid Computing*, pp.101–108. In proceedings.

Zhang, J. and Xiong, F. (2004) 'A novel watermarking for image security', Paper presented at the *19th International Symposium of Computer and Information Sciences (ISCIS 2004)*, Kemer-Antalya, Turkey. In proceedings.

Note:

[1]http://earth.google.com

[2]This figure is based on the time of this report.

[3]http://earth.google.com/coverage/coverage_list.pdf

[4]In the strictest sense, subjects include the users and the processes invoked by the users.

[5]Available at http://cimic.rutgers.edu/spatial

[6]http://www.njmeadowlands.gov

[7]For example: http://www.geocode.com/eagle.html-ssi

[8]The PKI and public certificate can be used to authenticate not only the user's identity but also the user's past scene (geotemporal) credentials.