

PROTECTING PRIVACY FROM CONTINUOUS HIGH-RESOLUTION SATELLITE SURVEILLANCE

Soon Ae Chun
and Vijayalakshmi Atluri

Abstract Privacy refers to controlling the dissemination and use of personal data, including information that is knowingly disclosed, as well as data that are unintentionally revealed as a byproduct of the use of information technologies. This paper argues that the high resolution geospatial images of our earth's surface, produced from the earth observing satellites, can make a person visually *exposed*, resulting in a technological invasion of personal privacy. We propose a suitable *authorization model for geospatial data* (GSAM) where controlled access can be specified based on the region covered by an image with privilege modes that include **view**, **zoom-in**, **overlay** and **identify**. We demonstrate how access control can be efficiently enforced using a spatial indexing structure, called MX-RS quadtree, a variant of the MX-CIF quadtree.

1. INTRODUCTION

In the year 2000 and the new millennium, there will be 31 satellites capable of providing land cover data at resolutions of 1 to 30 meters in orbit; 14 of these will be privately funded by US corporations, all with resolutions of 10 meters or better. As low-cost, highly responsive commercial satellite systems become operational, high resolution imagery is expected to become a regular input to consumer products and information services.¹ Remote sensing data sales and services are predicted to grow into a \$2 billion dollar market by the beginning of the 21st century [1].

There are numerous benefits to society in the constructive use of low cost satellite imagery. Examples include environmental monitoring, map making, disaster relief, infrastructure planning, national security, pinpointing of prospective sites to aid miners and drillers in planning access to natural resources, and detecting distressed crops early before such stress is visible to the human eye. Up-to-date satellite images can assist businesses in planning the placement of consumer outlets and manu-

facturing facilities, and help demographic analysts locate their target markets. Images can be used to aid police and fire crews to respond more quickly to distress calls, and to direct vehicle flows depending on observed traffic situations.

Motivation: While high resolution low cost satellite imagery enjoy many benefits, there are significant threats to privacy due to the commercial availability of high-resolution imagery in near real-time fashion. Public entities, such as local governments or public utility companies, collect, use and disseminate large amounts of personal information. Combination of this publicly available personal data pools with high resolution image data coupled with the integration and analysis capabilities of modern GIS systems providing geographic keys such as longitude and latitude, can result in a technological invasion of personal privacy. A person can, not only be identified by name or address, but can be *visually exposed*. Therefore, in the near future, it may be technically feasible for anyone to observe, record and measure the outdoor activities of anyone, at any place in the world (from backyard pools to nuclear plants), almost at any time. For example, one can clearly identify the objects in the high-resolution image shown in figure 1. Many scenarios can be envisioned that may threaten the privacy of individuals or organizations; some are listed below.

- Observation of military operations or movements of agents of foreign countries can be achieved by the click of a mouse [?].
- Unauthorized surveillance of a person's outdoor activities by a stalker or a burglar may help planning a break-in of a home. Tracking of residents entering and leaving the house through observing high resolution images over a period of time can simply be done on his computer.
- Tracking of the shipping volumes and patterns of a company by observing the number of trucks being loaded and unloaded can be valuable for a competing business enterprise.

These are some scenarios that depict the need for access control for high resolution geospatial image data. Although there are no policies or laws in place yet, they appear to be inevitable [?].

Aside from protecting privacy of individuals from near real-time high-resolution satellite surveillance, the need for controlled access to images arises because of different reasons:

- Concept based filtering: Filtering of images is needed, for example, to prevent children from accessing objectionable images available



Figure 1. A high resolution image

on the web. While traditionally access control is provided at the server, filtering requires access control at the client.

- **Controlled access to images:** Prevention of unauthorized access may be needed for providing controlled distribution of images to subscribers.
- **Content based access control:** Prevention of access may be needed for certain images based on their content, for example, to prevent public from accessing images of all vehicles with a color distribution used by the military.

Related Work: While there exist no work on providing access control for geospatial images, recently, a number efforts have been made to screen objectionable images using shape detection, object recognition, people recognition, face recognition, and content-based image retrieval. They include (1) filtering of images of naked people using a skin filter and a human figure grouper [3, 4], and (2) using a content-based feature vector indexing where an image is matched against a small number of feature vectors obtained from a training database [13, 14]. However, these approaches filter all images that match a set of criteria, but do not provide controlled access that facilitates access to images for legitimate users.

Our Contribution: A suitable access control for protecting privacy due to unauthorized high-resolution surveillance should not only be based on the spatial extent of images but also be based on their resolution. This is because, while a low resolution image may be revealed to the user regardless of its location coordinates, a high resolution image may not be

allowed access, except in the region where the user has access permission. For example, a factory owner may access every detail pertaining to his own operations, but should be prohibited from accessing the images that reveal the details of his competitor's operations. To the best of our knowledge, there does not exist any authorization model suitable for geospatial images. In this paper, we propose an authorization model that can provide access control for geospatial images based on their spatial extent and resolution, called *Geo-Spatial Authorization Model* (GSAM). Our access control model will use publicly available user information, such as property ownership and voter registration records to determine the spatial extent that the user is allowed to access, which in turn is used to determine the appropriate image(s), or a portion of it, from the image database.

To accomplish this, GSAM supports, in addition to the conventional privilege modes such as `read`, `insert`, `delete` and `modify`, privilege modes such as `view`, `zoom-in`, `overlay` and `identify` that can be defined based on the allowed resolution level for a given user/subject. To provide efficient processing of queries, images in a spatial database are typically stored using a spatial indexing structure. We have realized that an appropriate spatial indexing structure be based on image resolution to provide efficient processing of image operations, which will in turn effect the performance of evaluating an authorization in GSAM. Since there does not exist a suitable indexing structure, we first propose an enhancement to the spatial access method MX-CIF quadtree [?, ?], called MX-RS quadtree. Since MX-RS quadtree is capable of representing higher resolution images at lower levels of the tree, it facilitates efficient enforcement of privilege modes including the `zoom-in`, `view`, `overlay`, and `identify`, simply by controlling the traversal of the tree down towards higher resolution images.

We provide access control in two ways. (1) We *control the depth* a user can traverse, thereby controlling the resolution of the images (s)he can access. For example, anyone can access a low resolution image such as the New Jersey state map, but access to a 1 meter resolution image of an individual's house is prohibited as it may infringe on the privacy of that individual. (2) We *control the extent* a user can view. That is, a user is given access to high resolution images (say 1 meter), only for certain regions (typically the property (s)he owns, public parks, etc.) but not to all.

Organization of the paper: This paper is organized as follows. Section 2 provides a brief introduction to satellite and other geospatial image data as well as image databases. While Section 3 presents GSAM,

Section 4 shows how access control can be accomplished. Section ?? proposes an efficient access control approach that allows implementing it on top of the geospatial indexing structure. Finally Section 5 presents conclusions as well as future research we intend to pursue in this area.

2. BACKGROUND ON GEOSPATIAL IMAGES

Geospatial images can either be *digital raster images* that store image as a number of pixels, or *digital vector data* that store image as points, lines and polygons. Typically, satellite images, digital orthophoto quads and scanned maps are raster images, while maps of vector type (e.g. Shape file), digital line graphs, or census TIGER data are vector images. Other non-image geospatial data sets are data with locational information, such as census data, voter registration, land ownership data, and land use data.

Since the main focus of this paper concerns protecting privacy from high-resolution satellite surveillance, we provide more details on satellite imagery. Satellite images are a product of Remote Sensing. Remote sensing is a technology for sampling radiation and force fields to acquire and interpret geospatial data. Geospatial data are used to develop information about features, objects, and classes on Earth's land surface, oceans, and atmosphere. Remote sensing of the Earth traditionally has used reflected energy in the visible and infrared and emitted energy in the thermal infrared and microwave regions. It gathers radiation that can be analyzed numerically or used to generate images whose variations represent different intensities of photons associated with a range of wavelengths that are received at the sensor. Satellite images are pictorial representation of target objects and features in different spectral regions. Each of different sensor (commonly with bandpass filters) is tuned to accept and process the wave frequencies (wavelengths) that characterize each region. Each region normally shows significant differences in the distribution (patterns) of color or gray tones. A chief use of satellite image data has been in classifying different features in a scene into meaningful categories or classes. The image then becomes a thematic map (the theme is selectable, e.g., land use; geology; vegetation types; rainfall).

Satellite data have the following characteristics:

- The satellite's orbital information is changing; hence it is hard to obtain images whose spatial coverages are exactly the same.
- There are variabilities of images coming from different satellites and sensors, even if they observe the same region. Typically dif-

ferent sensors capture different characteristics of earth surface, e.g. land coverage and weather.

- Different sensors provide images of different resolution levels, from low to high. For example, the Advanced Very High Resolution Radiometer (AVHRR) is a broad-band, four or five channel (depending on the model) scanner, sensing the visible (red, green, blue), near-infrared, and thermal infrared portions of the electromagnetic spectrum. It produces 1km resolution images. Landsat Thematic Mapper (TM) provides multi-spectral imagery at 25m ground resolution. Radar sensors can transmit 5 to 10 meter resolution images. Sensors from the IKONOS satellite launched by Space Imaging/EOSAT promises to provide 1m Panchromatic and 4m Multispectral (blue, green, red, near-IR) data.
- For any remotely sensed image, there is a trade-off between spatial resolution, area of extent, and data volume. If the data volume is to be held constant, a high-resolution image will cover a small area, while a low-resolution image will cover a large area. The systems intended for the identification of land cover and land use have focused on moderate resolutions between 5 and 30 meters and swaths of 100 to 200 kilometers, while the high resolution satellites are designed with 1 to 3 meters resolution and 4 to 40 kilometer swaths.
- Each satellite image undergoes the process of georectification which involves two steps: georegistration and geocorrection. Geocorrection of the image is needed since the distances and directions in satellite images do not correspond to true distances and directions on the ground due to the variability of satellite position. Georegistration process registers each image with a known coordinate system (e.g. longitude, latitude), and reference units (e.g. degrees) and coordinates of left, right, top and bottom edges of the image.

In addition to these characteristics specific to remotely sensed data, satellite images also share characteristics of spatial data such as maps, since they cover spatial areas of the earth surface. In order to provide efficient retrieval of geospatial images based on spatial locations, and to support approximate matching (similarity based retrieval), range and nearest neighbor queries, various *spatial access methods* (SAMs) have been developed. They include approaches that transform rectangles into points in a space of higher dimensionality, e.g. grid files [8]; approaches that use linear quadtrees [12, 5], z-ordering [10] or other space filling curves [9]; and approaches based on trees (R-tree [7], MX quadtrees

[?, ?], k-d-trees [2], k-d-B trees [11], hB trees [?] and cell trees [6]). In this paper, we propose extension to the MX-CIF quadtree, called MX-RS quadtree, that enables efficient implementation of various privilege modes of GSAM.

3. AUTHORIZATION MODEL FOR GEOSPATIAL DATA (GSAM)

In this section, we formally present GSAM, an authorization model suitable for providing controlled access to geospatial data. Let $S = \{s_1, s_2 \dots\}$ denote a set of subjects, $O = \{o_1, o_2 \dots\}$ a set of objects, and $M = \{view, zoom-in \dots\}$ a finite set of privilege modes. In the following, we describe in detail the image objects and privilege modes, and present the formalism for authorization specification.

3.1. IMAGE OBJECTS

Image objects can either be raster or vector images. Vector objects describe geographic map features such as roads, parcels, soil units, or forest stands. It can contain several feature classes, such as arc, node, polygon, label point, annotation, tic, and coverage extent.

Each raster image object O_i is represented as a tuple, $\langle id, l, g, h, w, r, t \rangle$, where id is a unique identifier and l, g, h , and w are *latitude, longitude, height, width*, respectively, that represent the spatial extent of the image. r is for *resolution* of O_i , while t represents the *download timestamp*.

There is a set of access functions associated with each object. Given an image object, O_i , the function $rectangle(id)$ would retrieve the rectangular region (l, g, h, w) of the object. Similarly $resolution(id)$ would return r .

Each vector object, O_v , is represented as a tuple, $\langle id, l, g, h, w, t, k \rangle$, where id is a unique identifier and l, g, h , and w are *latitude, longitude, height, width*, respectively, that represent the spatial extent of the vector file. t represents the *last update timestamp*. k represents a *link* that links tabular data of geographic features contained in the vector object, O_v . Table 1 shows an example of tabular data records that can be linked to vector objects.

3.2. PRIVILEGE MODES

In our model, we support two types of privilege modes – *viewing* and *maintenance*. The viewing modes include *view, zoom-in, overlay*, and *identify*, and the maintenance modes are *insert, delete* and

Id	Area	Town	Acre	Landuse	Owner	Address
1	1216235.75	Moonachie	27.92000	vacant	RALPH ZISA	18 GARFIELD ST
2	7474.30	Carlstadt	0.17000	residential	WILLIAM LUBERTO	18 RIVERSIDE AVE
3	3077350.50	Teterboro	70.65000	commercial	JOAN VALENTINE	120 EGG HARBOR RD
4	44063.28	Little Ferry	1.01000	residential	HANNA Y. KARROUM	12 SCENIC VISTA DR
5	14740839.00	Teterboro	338.40000	office	ANTONIO TORRES	33 WASHINGTON AVE
6	363254.62	Ridgefield	8.34000	warehouse	LAWRENCE P. MEDICH	37 WASHINGTON AVE

Table 1. Tabular Data Linked to vector objects

update. The **view** privilege allows a user to see an image object covering a certain geographic area within a permitted resolution level.

The **zoom-in** privilege allows a user to view an image covering a certain geographic area at a higher resolution. Unlike conventional privilege modes that allow or deny access, this privilege specifies the level of zoom-in allowed, and is therefore expressed with an associated value, called **zoom level** (for example, **zoom-in:10**). The access control algorithm interprets this value and determines the level of resolution of the image that is allowed to be viewed by the user. Note that given an image, zooming-in can also be achieved using zoom-in algorithms, but the quality of the image decreases so that the result becomes useless, if zooming is done beyond a certain level. (For example, Figure 3 shows an image after the zoom-in operation to the rectangular region shown in Figure 2.) Thus the level of zoom-in a user is allowed should be determined based on the level (s)he can attain after applying the **zoom-in** algorithm. That is, if a user is allowed a zoom-in level of l_z , the access control algorithm must make sure that the user is given an image with a resolution of at most r that can not be zoomed-in to a resolution higher than l_z without losing its content. The functionality of providing the desired level of zoom-in is achieved by storing multiple images with different levels of resolution. Thus, if a user is allowed to access a region at a certain level of resolution, zooming-in is accomplished by retrieving a higher resolution image.

The **overlay** privilege allows users to generate composite images, where a composite image is constructed by *overlaying* one image on top of another. Although each individual image in isolation can be viewed by a user, sometimes an overlaid image may reveal more information than the user is allowed to access. For example, Figure 4 shows a 1m resolution aerial orthophoto image overlaid with a zoning vector coverage, represented in white lines. While a user can view both images individually, overlaying the street map on a high resolution image may help pin-pointing a person’s private property and viewing it in realtime.



Figure 2. Landsat Image with 28.5 meter resolution



Figure 3. Landsat Image with zoom-in level of 10

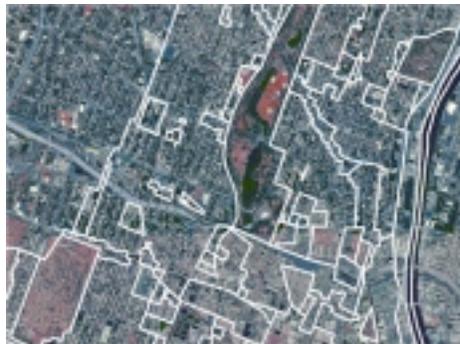


Figure 4. Orthophoto image overlaid with a zoning vector map

The **identify** privilege allows the user to view the tabular data linked to an image. As can be seen from Table 1, some of the data linked to the image, for example the ownership information, when shown with a high resolution image may provide visual exposure of a person's private property.

While the `insert` privilege allows a user to insert an image objects into the database, the `delete` allows her to remove images. The `update` privilege allows a user to replace one image with another as well as modify the attributes of the image, such as latitude, longitude, resolution, and link. In addition, it allows the user to update the tabular data linked to the image.

3.3. AUTHORIZATION

An authorization in GSAM is specified as follows:

Definition 1 An authorization a is a triple $\langle sub, obj, pr \rangle$, where

sub is a subject $s \in S$,

obj is (i) an object id of an object $o \in O$,

(ii) a region represented as a rectangle with (latitude, longitude, height, width), or

(iii) a set of object ids , and

pr is (i) a single privilege mode $m \in M$ or

(ii) a set of privilege modes $\{m_1, m_2, \dots\} \subseteq M$.

An object in our authorization specification can be a single image, a set of images, or a region. Although the region could be any polygon, for the sake of simplicity, in this paper, we limit it to represent only rectangles.

The privilege pr in an authorization triple may be composite, that is, may contain more than one privilege mode, which is especially useful when used with `overlay`. That is the case because, a subject may be allowed to overlay an image over another low resolution image, but not over a high resolution image. In order to specify such access control policies, we need a combination of both `zoom-in` and `overlay`.

In our model, as can be seen from the above definition, authorizations will allow one to specify that a subject is allowed to view a specific image or region with a specific resolution, or allowed to overlay a set of images with a specific resolution.

Following are some examples of authorizations.

$a_1 = \langle John, (50, 60, 10, 10), (zoom - in : 8) \rangle$

$a_2 = \langle Mary, 123, view \rangle$

$a_3 = \langle Ann, \{123, 456\}, overlay \rangle$

$a_4 = \langle Tom, \{123, 456\}, (overlay, *, 8) \rangle$

Above authorizations can be interpreted as follows: a_1 specifies that John is allowed to access a region centered at point (50,60) with width and height of 10, with a zoom-in level of 8. a_2 specifies that Mary can

view the object with the object id 123. a_3 specifies that Ann is allowed to overlay objects 123 and 456. Finally, a_4 specifies that Tom is allowed to overlay images 123 and 456 where the highest resolution level of object 456 is 8.

We use $a(sub)$, $a(obj)$ and $a(pr)$ to denote the subject, object and privilege of a , respectively. Moreover, to denote the attributes of each component in a , we use the notation $component_{attribute}$. For example, $a(pr_{zoomin})$ represents the zoom-in level specified in the privilege mode of a . We denote the set of all authorizations as *geo-spatial authorization base*, *GSAB*.

4. ACCESS CONTROL

When a subject requests to access images covering a specific geographic region at a specific resolution level, the access control mechanism must evaluate whether such request can be granted. We define the Access Request by a user, ur , as follows:

•

Figure 5. The System Architecture

Definition 2 [Access Request] An *access request* is a triple $ur = \langle s, o, pr \rangle$, where s is the subject, pr is the privilege mode, and o is the object which can be either of the following two: (i) a tuple (l, g, h, w, r) where (l, g, h, w) represents the requested rectangle represented with latitude, longitude, height and width, and r represents the level of resolution, and (ii) a set of object *ids*.

According to the above definition, a user may request to access an object by specifying its object id, or may request to access a rectangular region by specifying its latitude, longitude, height and width. We use $ur(s)$, $ur(o)$ and $ur(pr)$ to denote the subject, object and privilege mode specified in ur , respectively.

When a subject requests to access images covering a specific geographic region at a certain resolution level, the access control module (refer to figure 5) verifies whether there exists an authorization such that the object region specified in the authorization *overlaps* with (or contains) the requested object area. As a first step, it determines all the authorizations relevant to the access request. Access is denied if no relevant authorization exists. Then the authorization evaluation module determines either a set of object ids or a rectangular region that is allowed to be viewed by the subject and sends a request to the image database.

Since the region allowed to be viewed by the subject may not match exactly with the image(s) returned, the images returned from the image database need to be edited, namely assembled and/or cropped. This function is performed by the image processing module.

Given an authorization base $GSAB$, the following algorithm describes how an access request ur with view, zoom-in and overlay modes can be evaluated.

```

Algorithm [Authorization Evaluation] {
input:    $ur$ 

output:  set of images

begin
1.  Find   the set of authorizations  $A(ur)$  in  $GSAB$  such that
    for each  $a \in A(ur)$ 
         $(a(s) = ur(s)) \wedge (a(pr) = ur(pr))$ 
2.  if  $A(ur) = \emptyset$ 
    then return ('Access denied')
    else {
        while  $A(ur) \neq \emptyset$  {
for each  $a \in A(ur)$ 
case  $ur(pr) = \text{'view'}$ : {

                if  $((a(o) \text{ is } id) \wedge (ur(o) \text{ is } id))$ 
then
                    { if  $(a(o) = ur(o))$ 
                        then RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                    WHERE  $imageid = a(o)$  }

                    if  $((a(o) \text{ is } id) \wedge (ur(o) \text{ is not } id))$ 
                        then { if  $(\text{overlap}(\text{rectangle}(a(o)), \text{rectangle}(ur(o))) \neq$ 
                             $\emptyset)$ 
                                then RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                                    WHERE  $imageid = a(o)$  }
                                    if  $((a(o) \text{ is not } id) \wedge (ur(o) \text{ is not } id))$ 
                                        then {
                                             $area = \text{overlap}(\text{rectangle}(a(o)), \text{rectangle}(ur(o)))$ 
                                            RETRIEVE-IMAGES-WITH-AREA FROM ImageDB
                                            WHERE  $\text{overlap}(area(image), area) \neq \emptyset \wedge$ 
                                                 $resolution(image) \geq resolution(a(o))$ 
                                            PROCESS-IMAGES (area, images) }
                                         $A(ur) = A(ur) - a$  }
                                case  $ur(pr) = \text{'zoom-in'}$ : {
                                     $resolution(ur(o)) = ur(pr_{zoom-in})$ 
                                    if  $((a(o) \text{ is } id) \wedge (ur(o) \text{ is } id))$ 
                                        then { if  $(a(o) = ur(o))$ 
                                            then RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                                                WHERE  $imageid = a(o)$  }
                                            if  $((a(o) \text{ is } id) \wedge (ur(o) \text{ is not } id))$ 

```

```

    then { if ((resolution(a(o)) ≤ resolution(ur(o))) ∧
overlap(rectangle(a(o)), rectangle(ur(o))) ≠ ∅)
    then RETRIEVE-IMAGE-WITH-ID FROM ImageDB
        WHERE imageid = a(o) }
        if ((a(o) is not id) ∧ (ur(o) is not id))
        then { if (resolution(a(o)) ≤ resolution(ur(o)))
    then { area = overlap(rectangle(a(o)), rectangle(ur(o)))
RETRIEVE-IMAGES-WITH-AREA FROM ImageDB
WHERE (overlap(rectangle(image), area) ≠ ∅) ∧
        resolution(image) = resolution(ur(o))
}
        PROCESS-IMAGES (area, images) }
        A(ur) = A(ur) - a }
    case ur(pr) = 'overlay': {
    if ((a(oi) is id) ∧ (ur(oi) is id) ∧ ((a(oj) is id) ∧ (ur(oj) is
id)))
    then { RETRIEVE-IMAGE-WITH-ID FROM ImageDB
        WHERE imageid = a(oi) ∪ imageid = a(oj) }
        if ((a(oi) is id) ∧ (a(oj) is id) ∧ ((ur(oi) is not
id) ∧ (ur(oj) is not id)))
        then {
            if (overlap(rectangle(a(oi)), rectangle(ur(oi))) ≠
∅)
            then { RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                WHERE imageid = a(oi) ∧ resolution(image) ≥
resolution(a(oi)) }
            if (overlap(rectangle(a(oj)), rectangle(ur(oj))) ≠ ∅)
            then { RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                WHERE imageid = a(oj) ∧ resolution(image) ≥ resolution(a(oj)) }
                area = overlap(rectangle(ur(oi)), rectangle(ur(oj)))
                PROCESS-IMAGES (area, images) }
            if ((a(oi) is not id) ∧ (a(oj) is
not id)) ∧
                ((ur(oj) is not id) ∧ (ur(oj) is not id))
            then { Ri = overlap(rectangle(ur(oi)), rectangle(a(oi)))
                Rj = overlap(rectangle(ur(oj)), rectangle(a(oj)))
            if (overlap(Ri, Rj) ≠ ∅)
            then {
                RETRIEVE-IMAGES-WITH-AREA from IMAGEDB
                WHERE overlap(rectangle(image), Ri) ≠ ∅ ∧
resolution(image) ≥ resolution(a(oi))
resolution(image) ≤ resolution(ur(oi))
                RETRIEVE-IMAGES-WITH-AREA from
IMAGEDB
                WHERE overlap(rectangle(image), Rj) ≠
∅ ∧
resolution(image) ≥ resolution(a(oj))
resolution(image) ≤ resolution(ur(oj))
                PROCESS-IMAGES (overlap(Ri, Rj),
images) }}

```


tion model (GSAM). GSAM supports privilege modes including **view**, **zoom-in**, **overlay** and **identify** that are essential for providing constrained access to geospatial data based on the region covered by an image. We have demonstrated how access control can be efficiently enforced using a spatial indexing structure called MX-RS quadtree.

Our future research spans a number of directions. We intend to extend the authorization specification GSAM with temporal attributes. Unlike conventional authorizations that can be implemented as lists, authorizations in GSAM involve spatial attributes. In such a case, managing the authorization base and searching for authorizations based on the spatial extent is not trivial. Therefore, we intend to investigate techniques to maintain the authorization base. We plan to devise methodologies to verify the consistency of the authorization specification, analyze conflicts occurring due to simultaneous presence of contains, overlap and other operations, and strategies to resolve these conflicts. Another direction of future research is to build an indexing structure suitable for image access control in more general cases, where images at the same resolution level do not have fixed spatial extents. In addition, we intend to consider including the temporal aspects into the indexing structure. We also plan to investigate methods for providing refined access control where different geospatial information sets, such as health data and income data are integrated with image and map data.

Acknowledgments

The concept of access control for high resolution satellite imagery was conceived through discussions with Dr. Geoff Henebry, Biological Science Department, Rutgers University. We acknowledge Dr. Francisco Artigas, CIMIC, Rutgers University, for the information on geo-spatial images, their analysis and processing. The work was partially supported by the National Science Foundation under grant IRI-9624222 and the Meadowlands Environmental Research Institute as a grant from the Hackensack Meadowlands Development Commission.

Notes

1. The Ikonos 1 satellite system, launched on September 24, 1999, is capable of providing 1 meter resolution panchromatic images at 2 meter horizontal positional accuracy, provides a combination of very high resolution, high accuracy and high frequency revisit capability. <http://www.ssc.se/rst/rss/satinfo/ikonos.html>

References

- [1] Jonathan Ball. Satellite remote sensing. *TCS Remote Sensing and GIS web page*.

- [2] J Bentley. Multidimensional binary search trees used for associative searching. *CACM*, 18(9):509–517, September 1975.
- [3] M. Fleck, D. Forsyth, and C. Bregler. Finding naked people. In *Proceedings of 4th European Conference on Computer Vision*, pages 593–602, 1996.
- [4] D. et al Forsyth. Finding pictures of objects in large collections of images. In *Proceedings of International Workshop on Object Recognition*, pages 69 – 142, 1996.
- [5] I Gargantini. An effective way to represent quadtrees. *Comm. of ACM (CACM)*, 25(12):905–910, December 1982.
- [6] O. Gunther. The cell tree: an index for geometric data. Technical report, University of California, Berkeley, December 1986.
- [7] A. Guttman. R-trees: a dynamic index structure for spatial searching. In *Proceedings of ACM SIGMOD*, pages 47–57, June 1984.
- [8] K. Hinrichs and J. Nievergelt. The grid file: a data structure to support proximity queries on spatial objects. In *Proceedings of the WG'83 (International Workshop on Graph Theoretic Concepts in Computer Science)*, pages 100–113, 1983.
- [9] H. Jagadish. Linear clustering of objects with multiple attributes. In *Proceedings of ACM SIGMOD*, pages 332–342, May 1990.
- [10] J. Orenstein. Spatial query processing in an object-oriented database system. In *Proceedings of ACM SIGMOD*, pages 326–336, May 1986.
- [11] J. Robinson. k-d-b tree a search structure for large multidimensional dynamic indexes. In *Proceedings of ACM SIGMOD*, pages 10–18, April 1981.
- [12] Aref W. and H. Samet. Optimization strategies for spatial query processing. In *Proceedings of VLDB (Very Large Data Bases)*, pages 81–90, September 1991.
- [13] James Ze Wang, Jia Li, Gio Wiederhold, and Oscar Firschein. System for Classifying Objectionable Websites. In *Proceedings of the 5th International Workshop on Interactive Distributed Multimedia Systems and Telecommunication Services (IDMS '98)*, volume LNCS 1483, pages 113–124. Springer Verlag, September 1998.
- [14] James Ze Wang, Gio Wiederhold, and Oscar Firschein. System for Screening Objectionable Images Using Daubechies' Wavelets and Color Histograms. In *Proceedings of the 4th European Workshop on Interactive Distributed Multimedia Systems and Telecommunication Services (IDMS '97)*, volume LNCS 1309. Springer Verlag, September 1997.