CHAPTER 21
PROTECTING PRIVACY FROM CONTINUOUS HIGH-RESOLUTION SATELLITE SURVEILLANCE

Soon Ae Chun and Vijayalakshmi Atluri
MSIS Department and CIMIC
Rutgers University, Newark, NJ 07102
{soon,atluri}@cimic.rutgers.edu

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`.

1. INTRODUCTION

In the new millennium, 31 satellites, funded by both governments and private corporations, will be capable of providing land cover data at resolutions of 1 to 30 meters in orbit. 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. 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 manufacturing 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 enjoys 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 pool 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.

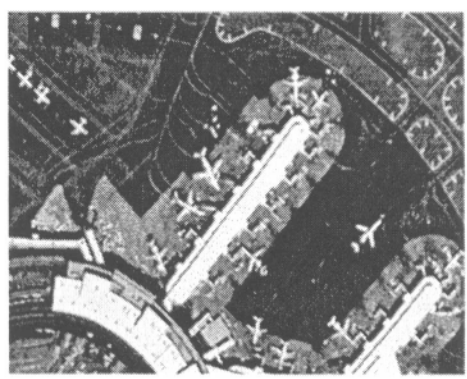

Figure 1. A high resoultion image

1. Observation of military operations or movements of agents of foreign countries can be achieved by the click of a mouse [7].

2. Unauthorized surveillance of a person's outdoor activities by a stalker or a burglar may help planning a break-in into 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.

3. 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 [7]. 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:

1. Concept based filtering: Filtering of images is needed, for example, to prevent children from accessing objectionable images available on the web. While traditionally access control is provided at the server, filtering requires access control at the client.

2. Controlled access to images: Prevention of unauthorized access may be needed for providing controlled distribution of images to subscribers.

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

Related Work: While there exists 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 [5, 6]. 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. While a low resolution image may be revealed to the user regardless of its location coordinates, a high resolution image may not be accessed, 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.

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 regions.

2. BACKGROUND ON GEOSPATIAL IMAGES

Geospatial images can either be *digital raster images* that store images as a number of pixels, or *digital vector data* that store images as points, lines and polygons. Typically, satellite images, digital orthophoto quads and scanned maps are raster images, while maps of vector type (e.g. a 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 regions 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 pic-

torial representation of target objects and features in different spectral regions. Each 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:

1. The satellite's orbital information is changing; hence it is hard to obtain images whose spatial coverages are exactly the same.

2. There are variabilities of images coming from different satellites and sensors, even if they observe the same region. Typically different sensors capture different characteristics of earth surface, e.g. land coverage and weather.

3. 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 electro-magnetic spectrum. It produces 1km resolution images. Landsat Thematic Mapper (TM) provides multispectral 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.

4. 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.

5. 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 registers each image with a known coordinate system (e.g. longitude, latitude), reference units (e.g. degrees) and coordinates of left, right, top and bottom edges of the image.

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 $S = \{s_1, s_2 \ldots\}$ denote a set of subjects, $O = \{o_1, o_2 \ldots\}$ a set of objects, and $M = \{view, zoom\text{-}in \ldots\}$ 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, $< id, l, g, h, w, r, t >$, where *id* is a unique identifier and *l*, *g*, *h*, and *w* are *latitude*, *longitude*, *height*, and *width*, respectively, that represent the spatial extent of the image. *r* is for *resolution* of O_i, while *t* represents the *download timestamp.* Each vector object, O_v, is represented as a tuple, $< id, l, g, h, w, t, k >$, where *id* is a unique identifier and *l*, *g*, *h*, and *w* are *latitude, longitude*, *height*, and *width*, respectively, that represent the spatial extent of the vector file. The symbol *t* denotes the *last update timestamp.* The symbol *k* denotes a *link* that links tabular data of geographic features contained in the vector object, O_v.

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*.

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 `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. 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 overlayed image may reveal more information than the user is allowed to access. Overlaying the street map on a high resolution image may help pin-pointing a person's private property and viewing it in realtime.

The `identify` privilege allows the user to view the tabular data linked to an image. 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 object into the database, the `delete` privilege 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 $\{m_1, m_2, \ldots\} \subseteq M$.

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 is 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 a request can be granted. We define the Access Request by a user, *ur*, as follows:

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 that consists of latitude, longitude, height and width, and r represents the level of resolution, or (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

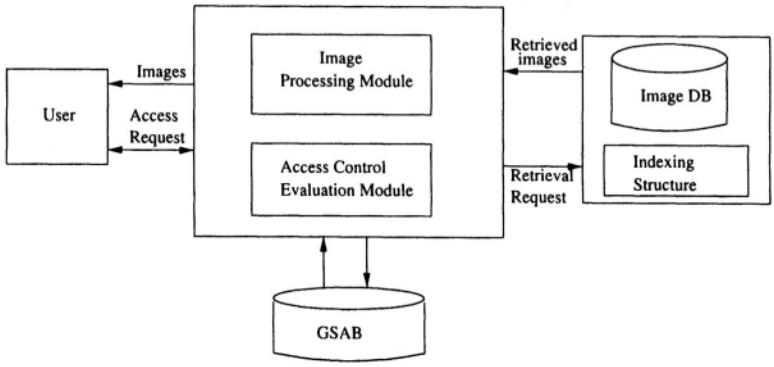

Figure 2. The System Architecture

$ur(s)$, $ur(o)$ 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 2) 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 1 [Authorization Evaluation]
input: ur
output: set of images
begin
1.Find the set of authorizations A(ur) in $GSAB$ such that
  foreach $a \in A(ur)$
    $(a(s) = ur(s)) \wedge (a(pr) = ur(pr))$
2.if$A(ur) = \emptyset$
then return (''Access denied'')
else {
  while $A(ur) \neq \emptyset$ {
    foreach $a \in A(ur)$
    case $ur(pr)$ ='view': {
      if $((a(o)$ is $id) \wedge (ur(o)$ is $id))$
      then{**if** $(a(o) = ur(o))$
        then RETRIEVE-IMAGE-WITH-ID FROM ImageDB
          WHERE $imageid = a(o)$ }
      if $((a(o)$ is $id) \wedge (ur(o)$ is not $id))$
      then{**if** $(overlap(retangle(a(o)), rectangle(ur(o))) \neq \emptyset)$

```
            then RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                 WHERE imageid = a(o) }
         if ((a(o) is not id) ∧ (ur(o) is not id))
         then {
            area = overlap(rectangle(a(o)), rectangle(ur(o)))
            RETRIEVE-IMAGES-WITH-AREA FROM ImageDB
            WHERE overlap(area(image), area) ≠ ∅ ∧
               resolution(image) ≥ resolution(a(o))
            PROCESS-IMAGES (area, images) }
         A(ur) = A(ur) − a }
      case ur(pr) ='zoom-in': {
         resolution(ur(o)) = ur(pr_zoom−in)
         if ((a(o) is id) ∧ (ur(o) is id))
         then{if (a(o) = ur(o))
            then RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                 WHERE imageid = a(o) }
         if ((a(o) is id) ∧ (ur(o) 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(o_i) is id) ∧ (ur(o_i) is id)) ∧ ((a(o_j) is id) ∧ (ur(o_j) is id)))
         then{ RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                 WHERE imageid = a(o_i) ∪ imageid = a(o_j) }
         if ((a(o_i) is id) ∧ (a(o_j) is id)) ∧ ((ur(o_i) is not id) ∧
              (ur(o_j) is not id)))
         then {
            if (overlap(rectangle(a(o_i)), rectangle(ur(o_i))) ≠ ∅)
            then { RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                 WHERE imageid = a(o_i) ∧ resolution(image) ≥ resolution(a(o_i)) }
            if (overlap(rectangle(a(o_j), rectangle(ur(o_j))) ≠ ∅)
            then { RETRIEVE-IMAGE-WITH-ID FROM ImageDB
                 WHERE imageid = a(o_j) ∧ resolution(image) ≥ resolution(a(o_j)) }
            area = overlap(rectangle(ur(o_i)), rectangle(ur(o_j)))
            PROCESS-IMAGES(area, images) }
         if ((a(o_i) is not id) ∧ (a(o_j) is not id))
              ∧ ((ur(o_j) is not id) ∧ (ur(o_j) is not id)))
         then { R_i = overlap(rectangle(ur(o_i)), rectangle(a(o_i)))
                R_j = overlap(rectangle(ur(o_j)), rectangle(a(o_j)))
                if (overlap(R_i, R_j) ≠ ∅)
                then {
                   RETRIEVE-IMAGES-WITH-AREA from IMAGEDB
                   WHERE overlap(rectangle(image), R_i) ≠ ∅∧
                      resolution(image) ≥ resolution(a(o_i))
                      resolution(image) ≤ resolution(ur(o_i))
                   RETRIEVE-IMAGES-WITH-AREA from IMAGEDB
                   WHERE overlap(rectangle(image), R_j) ≠ ∅∧
                      resolution(image) ≥ resolution(a(o_j))
                      resolution(image) ≤ resolution(ur(o_j))
                   PROCESS-IMAGES (overlap(R_i, R_j), images) }}
            A(ur) = A(ur) − a } } }
end

Procedure PROCESS-IMAGES
input: area, retrieved-images
output: images covering only area
begin
      foreach image i ∈ images
         chop (area, i)
      for each imageset I ∈ same resolution level {
         images = assemble-area (area, I)
```

```
        return(images) }
end
```

This algorithm considers three cases for evaluating each privilege mode. In the first case, both the access request and authorization are specified with image ids. In this case, evaluation of an access request is done by testing whether the ids are the same. In the second case, the access request is specified as a rectangular region, but the authorization is specified with an image id. In this case, evaluation involves determining the overlapping region of the image specified in the authorization with the requested region. If the overlapping region is empty, access is denied. Otherwise, appropriate request is sent to the image database to retrieve the image. The case where authorization is specified with a region and the access request is specified as an id can be dealt with in a similar manner. Therefore, this is not included in the algorithm. In the third case, both the access request and the authorization are specified as rectangular regions. In this case, the overlapped region must be determined first. The area is then used to retrieve the relevant images.

Further processing is done by the procedure PROCESS-IMAGES if the area covered by the retrieved images does not coincide with the region authorized to be viewed by the subject. In this case the image is cropped. If more than one image are retrieved, they are first assembled together before cropping.

5. CONCLUSIONS AND FUTURE RESEARCH

In this paper, we have argued that near-continuous surveillance through high resolution satellite images when combined with geographic information could be a threat to privacy. In order to address this issue, we presented a suitable access control model, called Geospatial Authorization 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. Our future research spans a number of directions. We plan 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. We have demonstrated in [2] how access control can be efficiently enforced using a spatial indexing structure called MX-RS quadtree. In future research

we will 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 Geoff Henebry. We acknowledge Francisco Artigas for the information on geo-spatial images, their analysis and processing. We thank James Geller for commenting on an earlier draft of this paper. 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 Hackenack Meadowlands Development Commission.

References

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

[2] Soon Ae Chun and Vijayalakshmi Atluri. Protecting privacy from continuous high-resolution satellite surveillance. Technical report, CIMIC, Rutgers University, November 1999.

[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] 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 Telcommunication Services (IDMS '98)*, volume LNCS 1483, pages 113–124. Springer Verlag, September 1998.

[6] 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 Telcommunication Services (IDMS '97)*, volume LNCS 1309. Springer Verlag, September 1997.

[7] Robert Wright. Private Eyes. *The New York Times Magazine*, September 1999.