Working Paper Series
ISSN 1177-777X

Location-based indexing for mobile context-aware access to a digital library

Annika Hinze and Wendy Osborn

Working Paper: 04/2007
August 22, 2007

Department of Computer Science
The University of Waikato
Private Bag 3105
Hamilton, New Zealand

Location-based indexing for mobile context-aware access to a digital library

Annika Hinze[1] and Wendy Osborn[2]

[1]The University of Waikato, New Zealand;
[2]The University of Lethbridge, Canada
hinze@cs.waikato.ac.nz, wendy.osborn@uleth.ca

Abstract

Mobile information systems need to collaborate with each other to provide seamless information access to the user. Information about the user and their context provides the points of contact between the systems. Location is the most basic user context.

TIP is a mobile tourist information system that provides location-based access to documents in the digital library Greenstone. This paper identifies the challenges for providing efficient access to location-based information using the various access modes a tourist requires on their travels. We discuss our extended 2DR-tree approach to meet these challenges.

1 Introduction

The TIP/Greenstone bridge [7] provides access to a set of rich materials with cross references between different digital library collections in Greenstone[13] and the tourist information system TIP [6]. The intention of the collaboration of the two information systems is that a user, traveling with an internet-connected mobile devise (e.g, a pocketPC) , is actively presented information based on their location and their interests; when a passing detail piques their interest, the user seamlessly taps into the "deeper" resources managed by the digital library that can better satisfy their quest for both more details and related information.

Another link between TIP and Greenstone has been forged in the Audio Book Service. Similarly, users access resources in the digital library (here: audio book files and book chapters) based on their current location and preferences.

This working paper explores the issues for *efficient* access to location-based information in such combined information systems. Section 2 provides and introduction to the diverse system components involved: the TIP system, the Greenstone system, the TIP/Greenstone Bridge, and the TIP/Greenstone Audio Books service. Section 3 explores the different query types of location-based

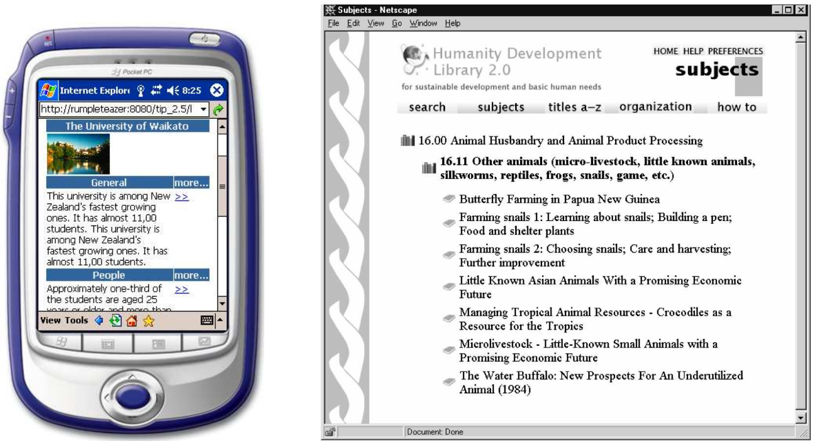

(a) mobile TIP (b) Greenstone: digital library system

Figure 1: Two main components of the hybrid system

access that are required from the user's point of view and the system perspective, and the challenges for index-based access that are created by these query types. Section 4 details background information about existing location-based indexing strategies. In Section 5, we introduce and explain our approach to efficient index-based access to location-dependent data.

This paper serves as a starting point for research into location-based indexing for mobile and context-aware access to digital libraries as well as other information systems. The aim of the paper is the identification of the issues that have to be addressed and the definition of our research strategy.

2 Background

This section describes the foundations of TIP and Greenstone as well as their collaborative extensions. We refer to the following publications for further details [5, 6, 7, 13].

2.1 Tourist Information Provider (TIP)

The Tourist Information Provider (TIP) system delivers location-based information to mobile users. The presented information is based on the user's context, such as their current location, their interest in particular semantic groups of sights and topics, and their travel history. Semantic groups and topics are captured in a user's profile. Examples for groups of sights are *public art*, *buildings*, and *beaches*; topics may be *history* or *architecture*. The travel history of a user includes the locations/sights that the user visited and the user's feedback about these sights.

Figure 1(a) shows the TIP standard interface in a mobile emulator. The user is at the University of Waikato. Their profile is *groups*={*buildings; parks*}; *topics* = {*architecture; history*}. The university is displayed as a building close to the user's current position. In addition to the core functionality, TIP supports several services such as recommendations and travel navigation on maps (for details see [6]). The TIP system combines an event-based infrastructure and location-based service for dynamic information delivery. The heart of the system is a filter engine cooperating with a location engine. The filter engine selects the appropriate information from the different source databases based on the user and sight context. Changes in the user's location are transmitted to the TIP server, where they are treated as events that have to be filtered. For the filtering, the sight context and the user context are taken into account. The location engine provides geo-spatial functions, such as geo-coding, reverse geo-coding, and proximity search. For places that are currently of interest, the system delivers sight-related information.

2.2 Greenstone Digital Library

Greenstone is a versatile open source digital library toolkit [13]. Countless digital libraries have been formed with it since its release on SourceForge in 2000: from historic newspapers to books on humanitarian aid; from eclectic multimedia content on pop-artists to curated first editions of works by Chopin; from scientific intuitional repositories to personal collections of photos and other document formats. All manner of topics are covered—the black abolitionist movement, bridge construction, flora and fauna, the history of the Indian working class, medical artwork, and shipping statistics are just a random selection. All manner of document formats are covered, including: HTML, PDF, Word, PowerPoint, and Excel; MARC, Refer, Dublin Core, LOM (Learning Object Metadata) and BibTeX metadata formats; as well as a variety of image, audio, and video formats. It also supports numerous standards including OAI-PMH (Open Archives Initiative Protocol for Metadata Harvesting), Z39.50 and METS to assist interoperability. See *www.greenstone.org* for more details.

Greenstone provides means to transform a set of digital resources into an organized, manageable collection accessible through browsing structures as well as direct access through fielded and full-text searching. The end result is shaped by a configuration process via the Greenstone Librarian Interface. Meta-data for documents and collections are used for shaping the collection building process and the final storage, access and presentation. This aspect of Greenstone is exploited—seeded through a gazetteer—for identifying and marking up place names.

In addition to the production-level version of the software used for the cited examples above (known as Greenstone 2), for exploratory purposes a digital library framework that utilises web services is available for research-based work. Called Greenstone 3 and backwards compatible, this forms the digital library

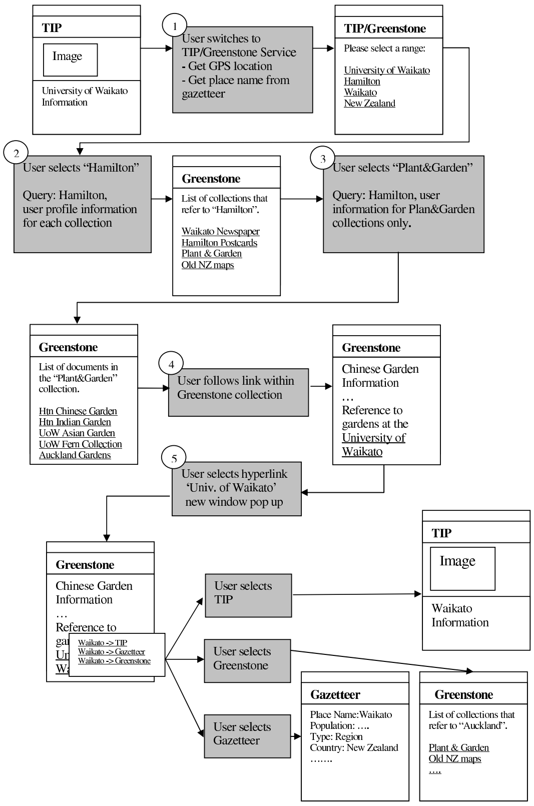

Figure 2: Overview of example interaction with TIP/Greenstone Bridge [7]

code base used for the hybrid system. In particular it allows for the fine-grained interaction necessary for the collaboration with TIP.

2.3 TIP/Greenstone Bridge

The TIP/Greenstone bridge is first known combination of a mobile tourist guide with a digital library.

An overview of the interactions is given in Figure 2. Starting from a TIP information page about the University of Waikato, the user decides to look up the digital library collections that refer to her current location. When she switches to the page of TIP/Greenstone Service, the system will display regions and places that are near-by that she might want to search for in the collection repository provided by Greenstone. This is necessary as her location is the

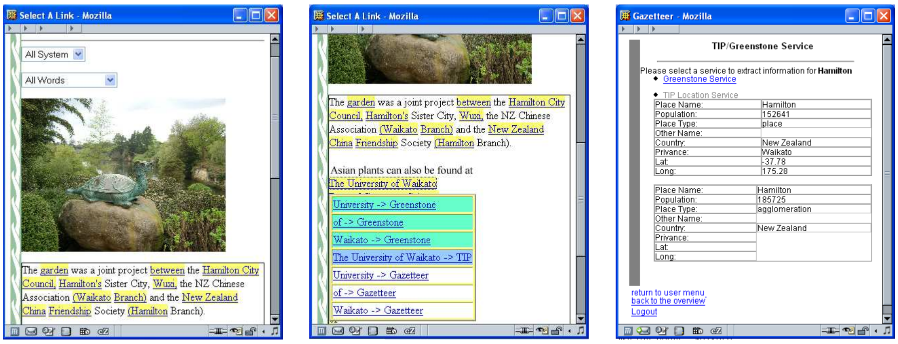

(a) highlighted text (context world) (b) back link pop-up (context world) (c) gazetteer (context New Zealand)

Figure 3: Overview of example interaction with TIP/Greenstone [7]

University of Waikato, close to the river Waikato, in the region Waikato, in New Zealand, on the north island, etc. All these locations could be used to search the library and the user can guide the selection. Based on the user's selection, the system triggers a location-based search in DL collections. The user is presented with a list of all collections that refer to the selected region.

After selecting the region 'Hamilton' in Step 2, the user has the choice between various DL collections, from which she chooses a collection that contains, among others, references to the local Hamilton Gardens (Step 3). The user selects a document about the Chinese Garden in Hamilton (Step 4). Figure 3(a) shows the Greenstone interface with the Chinese Garden Document. The user chooses to indicate all place names in the document: a special feature of the TIP/Greenstone service. All words that are recognised as place names are highlighted. Step 5 describes how the user can lookup the highlighted places in TIP or in Greenstone, via pull-down menu (see Figure 3(b)). The menu displays links only to existing data pages/documents. Different background colors indicate different target programs.

Figure 4 shows the conceptual interplay between TIP and Greenstone as they are connected via the Bridge, using the Gazetteer. A gazetteer provides information about locations, such as population, province, country (see Figure 3(c)). The lower part of Figure 4 shows the data that are retrieved.

For location-based access to Greenstone documents, the digital library documents have to be pre-processed and the locations identified within the document content. Our current implementation of the TIP/Greenstone bridge uses a Gazetteer to identify place names annotated by country. A simple location-aware mark-up of the documents is used in the software version described in [7].

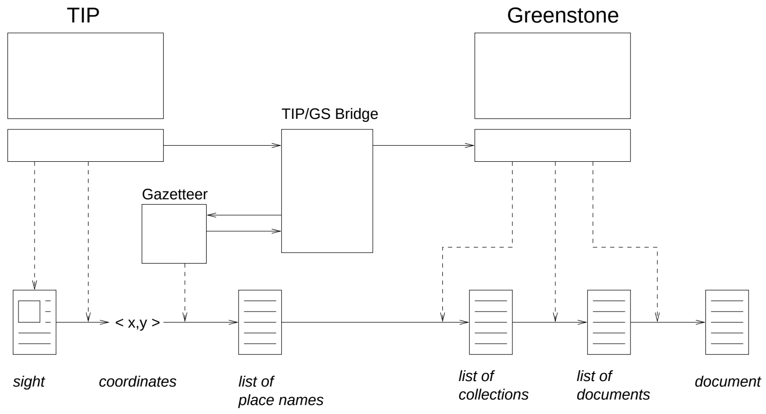

Figure 4: TIP/Greenstone bridge data flow

2.4 TIP/Greenstone Audio Books

The Greenstone DL can store a wide variety of data formats (see Section 2.2), including text and audio files. Two methods of audio books are supported in the TIP/GS Audio book service: audio books as readings from text documents and audio books provided as MP3 files. Detailed description of the audio book service can be found in [5]. In a similar way to finding text documents related to given locations (see Section 2.3), the audio book service accesses documents and mp3 files containing (interlinked) book chapters that are connected to selected locations.

The Audio Books service interactions are described in Figure 5. The service provides a list of books corresponding to the user's current location (Step 1). All related Audio Books will be listed (Step 2, Figure 5) with a general introduction (Step 3). For each book, the user can select available chapters from a chapter list (Step 4).

The Audio Books are categorized into reference books and general literature. The select-chapter function is only available for literature books, because the chapters of a reference book do not need to be read in a certain order. When users select the option to play the chapter via audio, the system will open the audio control panel. Users can also select to show the Audio Books text in Greenstone. An additional function of a Chapter Map is available for literature books: The Chapter Map offers the user to play chapters in sequence based on the user's travel route. The system lists all of the routes related to the current user with their sights and the related chapters. Chapters are read automatically when approaching a sight ion the selected route.

Similar to document access in Section 2.3, the Audio Book service provides location-based access to documents and mp3 files. Here the location-based access is directed by the meta-data defined for the items, not by analysing their content.

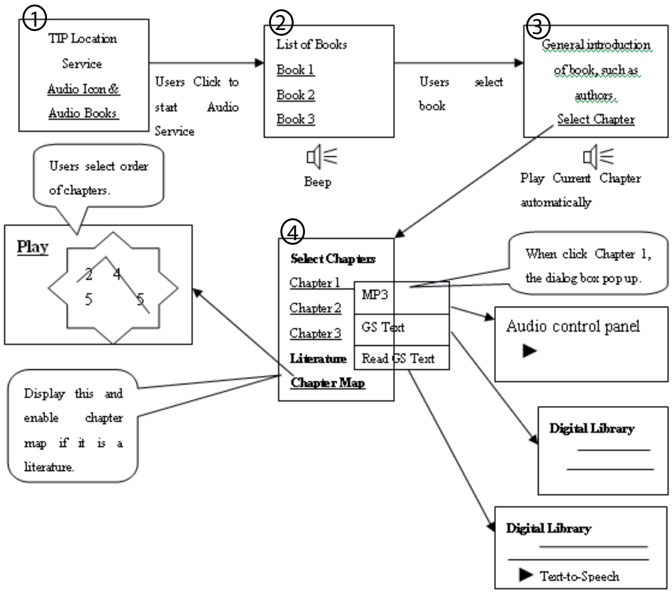

Figure 5: Overview of example interaction with TIP/Greenstone Audio Book service [5]

3 Location-based Access and Query Types

From a user's point of view, several types of inquiries can be started. We first describe the location-based retrieval from the perspective of the traveller and then analyze which types of queries are necessary to support the user.

3.1 Location-based Retrieval: User's Perspective

User at location
The user is situated at a location; we consider here them to be static regarding their location as the queries are concerned. At any new location they may issued another query. We can distinguish three types of requests with their respective responses

- A: *Documents about this sight?*
  Starting from a given place/sight at which the user is located or which they are browsing in TIP, the user wants to find related documents or collections within Greenstone. The input may be a point, as in the example in Figure 6(a), or a 2D object as in Figure 6(b).
  –input: TIP sight,
  –output: documents and collections

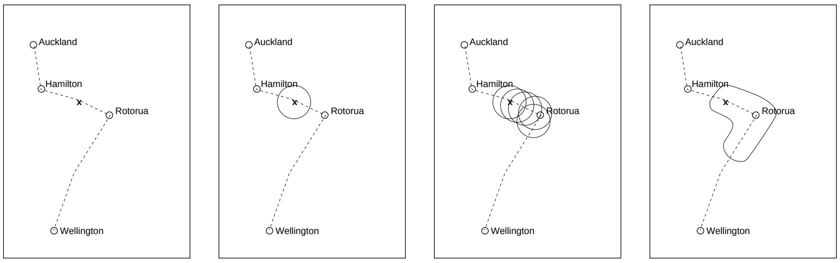

(a) Point queries

(b) Area queries

(c) Incremental 2D range queries

(d) 2D range queries

Figure 6: Examples of query types

- B: *Documents related to this place?*
  Within TIP starting from a certain coordinate the user is at (via GPS or selected on a map) they want to find related documents or collections within Greenstone. This is always starting from a single coordinate as in the example in Figure 6(a). Note the issue of layering and hierarchy: given the coordinates `<x,y>`, a number of spatial objects and concepts may be identified. An example is shown in Figure 7. Some of the objects belong to a hierarchy, such as the library building being part of the University of Waikato. Some objects are layered, as they belong to different concepts; for example the North Island as geo-spatial territory shares a subspace with New Zealand as the political concept of a country (see Figure 7).
  –input: Coordinates,
  –output: documents and collections

- C: *Anything related to the place mentioned?*
  This is a back-link function where out of the documents in Greenstone the use can then find new place names in the document and with those jump back to TIP, the Gazetteer, or Greenstone.
  –input: Place name in Greenstone,
  –output: links to TIP (place), Gazetteer (place), or Greenstone (place) documents and collections

In addition to the queries listed above there are the 'typical' location-based queries in TIP: based on the current coordinates of the the system find sights near-by. These queries are currently answered by TIP's postGIS database.

User moving – Changing locations
The user moves along a certain trajectory and continuously queries the system for information. An example is shown in Figure 6(c). In these queries the user

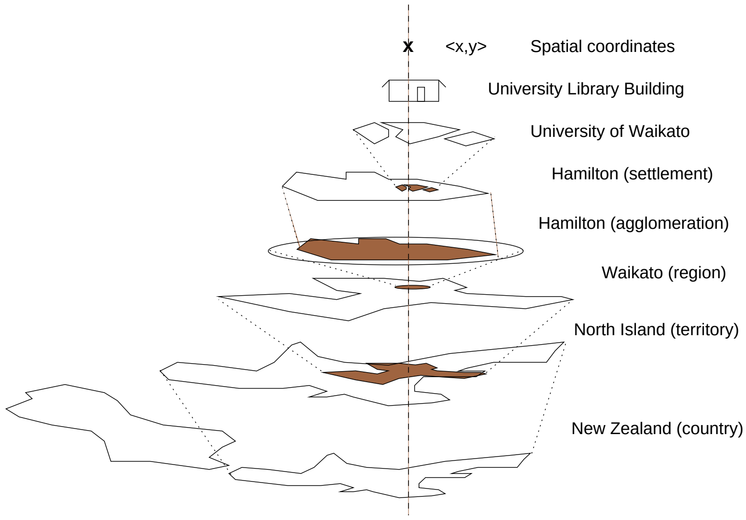

Figure 7: Hierarchy and Layering of geo-spatial and conceptual objects

starts at a given point and then gradually moves forward along a trajectory that is not known beforehand. We identify two types queries:

- A: *input: coordinates (near by starting-coordinates),*
  output: sights in TIP

- B: *input: coordinates (near by starting-coordinates),*
  output: documents and collection in Greenstone referring to places near by to this location (within ϵ range)

Virtual user travel

These are queries that are triggered, for example, by travel planning on the map. The user draws a line and wants to see everything that is close by, such as in teh example in Figure 6(d). The trajectory is know at the time of querying. Note the similarities to queries for areas work; for example, all documents/sights within New Zealand or within the Waikato region (see Figure 6(b)). We identify two types of queries:

- A: *input: coordinates along a line,*
  output: coordinates and IDs of sights in TIP

- B: *input: coordinates along a line,*
  output: documents and collection in Greenstone referring to places within the range of the locations

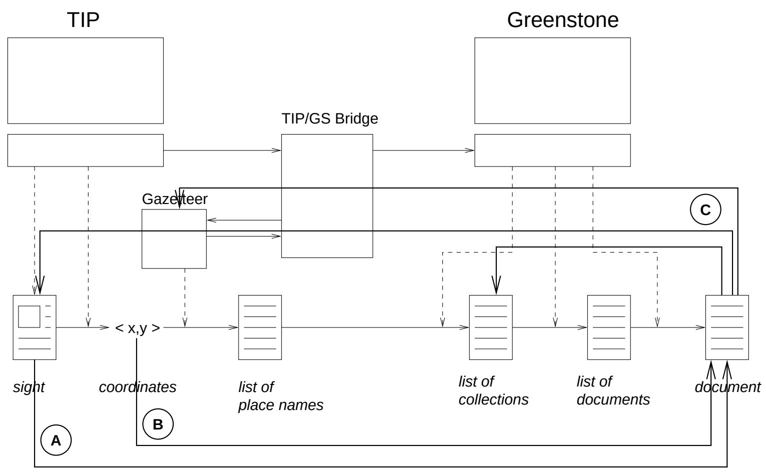

Figure 8: Point/range queries from and to documents

3.2 Location-based Retrieval: System's Perspective

The following queries need to be supported:

1. Point queries
   Point queries start from a single coordinate and query for items (documents, sights, files) that are associated with this location. Queries in TIP and in the TIP/Greestone bridge can be seen as point queries; for illustration see Figure 8.

   Point queries are not very typical as they refer to single `<x,y>` coordinates. The granularity of GPS data is too fine for items to be exactly at given coordinates. In addition, travelers typically are close to objects of interest but not standing directly on top of or in the object (where GPS may not work). Most point queries will therefore be enhanced by an ϵ radius to cover objects in close proximity to the given location. Consequently, the queries are executed as 2D range queries.

2. 2D range queries
   As argued before, most point queries are executed as 2D range queries. In addition, genuine range queries are triggered, for example, by travel planning on the map (c.f. virtual user travel). Range queries refer to polygon shapes and retrieve all objects linked to points within the polygon; area queries are types of 2D range queries. 2D range queries are typical queries while a user is immobile. Efficient access for 2D range queries requires location-based indexes.

3. incremental 2D range queries
   Incremental range queries follow a trajectory that is evolving. The are

issued by mobile users. They can also be viewed as range queries starting from a near-by point; see Figure 6(c). Efficient access for incremental queries requires preservation of information of neighborhood information between objects within the index.

3.3 Issues identified

From the analysis presented above, we identify the following challenges that need to be addressed for the efficient support of location-based mobile information access:

1. Efficient execution of 2D range queries and incremental range queries. Index structures need to be developed that provide direct and efficient support for the identified query types. Efficiency is especially important in a mobile environment where the user's location may be constantly changing.

2. Support for a wider notion of context, such as time or user preferences. As argued before, information access depends on a variety of context data. Indexes that ignore the other context dimensions will encounter mapping problems.

3. Index structures need to support a distributed setting where location-based data may be provided by various providers, depending on the location and available connectivity of the mobile users.

4. Location-based indexing for mobile applications need to explicitly support hand-held mobile devices. In addition to efficient retrieval, location-based and context-aware caching on devices with limited capacity or support for caching is also an important issue.

The next section introduces background on spatial indexing as foundation that will allow is to discuss new index structures in the subsequent section.

4 Background in Spatial Indexing and Retrieval

This section introduces background information on spatial data, typical data structures, access methods and index structures. A more extensive introduction can be found in [4, 8, 12]

4.1 Multidimensional (Spatial) Data

Spatial data is any type of non-standard data that exists in multidimensional space. It ranges in complexity from simple points to objects composed of sub-objects. All objects are composed from the following basic types [11, 12]:

- *Point.* A point is a zero-dimensional object that represents a particular location in n-dimensional space. For example, a town is represented on a map using a point.

- *Lines and Line-strings.* A line exists between two points and represents a one-dimensional feature in n-dimensional space. A line-string represents a sequence of one-dimensional features. For example, a line represents a road segment, while a line-string represents an entire road.

- *Polygon.* A polygon represents a two-dimensional feature in n-dimensional space. For example, a polygon is used to represent larger municipal areas. It can be arbitrarily shaped. Many multidimensional structures use an approximation to represent polygons in the form of a minimum bounding rectangle.

Using these basic types, more complex objects can be formed. For example, a province consists of many towns (i.e. points), cities (i.e. regions), and roads (i.e. linestrings).

For example, the region of Waikato contains towns (i.e. points), cities (i.e. regions), and roads (multiple connected lines, or linestrings). Each feature has a location in space. In addition, all features are spatially related to each other (for example, the city of Rotorua is located southeast from the city of Hamilton)

4.2 Search in Multidimensional Space

Two important issues for large spatial data sets are the efficient retrieval of a specific object and the efficient search for subsets of spatial objects. Some common search types include [2, 4, 11, 12]:

- *Exact Match.* The exact match search finds a matching point, object or object approximation.

- *Partial Match.* The partial match search locates one or more points that match a subset of values for a query point.

- *Point.* The point search finds all points or objects that overlap a query point. Examples for our system are point queries that identify objects higher up in the hierarchy, such as towns or countries (see Figure 7).

- *Region.* The region search finds all objects and points that overlap a search region. It is assumed that if an object overlaps the search region on its boundary only, then it is considered overlapping. This search region is usually defined as a k-dimensional "rectangle", which is supported by most data structures. These are typical static searches for places in TIP and Greenstone.

- *Containment.* The containment search finds all points and objects that are entirely contained in a query region.

- *Spatial Join.* The spatial join takes two data sets, X and Y, and finds all (x,y) pairs that satisfy some criteria (e.g. x overlaps y, x left of y, x right of y, *etc.*).
- *Nearest Neighbour.* The nearest neighbour search finds k closest objects or points to a query object or point. Examples are searches for sights near by in TIP.

4.3 Multidimensional Access Methods

The B-tree [1] and its variants provides efficient access for atomic (i.e. one-dimensional) values but they are not suitable for data that are described in multiple dimensions. Günther and Gaede [4] state that no n-dimensional to one-dimensional mapping of spatial data exists that preserves all spatial relationships. Therefore, a solution to the spatial relationship preservation problem that maintains the existing properties of spatial access methods is to use a hierarchical structure that fits the n-dimensional data as given instead of forcing the data to fit the structure.

Many multidimensional access methods are proposed in the literature and can be classified into three categories [4]. The first are main memory access methods that extend the binary search tree (BST) to handle multidimensional points. The second are point access methods that consider tree balance and secondary storage management issues when handling point data. The third are spatial access methods that extend the B-tree and their variants to uniformly store both point and object data.

4.4 The 2DR-tree

The nDR-tree uses nodes that are the same dimensionality as the data. The two-dimensional version, the 2DR-tree uses two-dimensional nodes to index two-dimensional data: The minimum bounding rectangles in each node are organized according to a "validity rule" that preserves spatial relationships. This provides support for both binary searching that takes advantage of spatial relationships and greedy searching that reduces the number of minimum bounding rectangles (MBR) within a node that must be tested.

The spatial relationships supported in the 2DR-tree are north, northeast, east, southeast, south, southwest, west and northwest. A spatial relationship is defined between two objects using the centroids of their MBRs. The centroid of an MBR are the co-ordinates (i, j) of its centre. For example, MBR 1 is northeast of MBR 2 if the centroid of MBR 1 is northeast of the centroid of MBR 2.

4.4.1 Node Validity

To employ different binary searching strategies, the spatial relationships between MBRs in each node must be preserved. For each location $N_{(i,j)}, i = 0 \dots (X - 1), j = 0 \dots (Y - 1)$ in node N, if $N_{(i,j)}$ contains an MBR $MBR_{(i,j)}$,

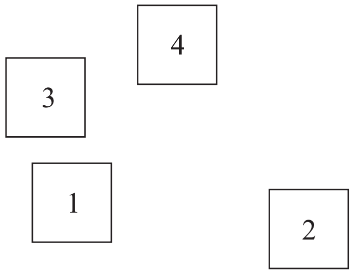

(a) Set of MBRs

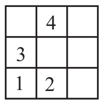

(b) Valid

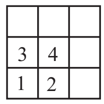

(c) Invalid

Figure 9: Node Validity

- Location $N_{(k,l)}, k = (i+1)\dots(X-1), l = 0\dots j$ contains $MBR_{(k,l)}$ whose centroid is south, east or southeast of the centroid for $MBR_{(i,j)}$,
- Location $N_{(k,l)}, k = (i+1)\dots(X-1), l = (j+1)\dots(Y-1)$ contains $MBR_{(k,l)}$ whose centroid is northeast of the centroid for $MBR_{(i,j)}$, and
- Location $N_{(k,l)}, k = 0\dots i, l = (j+1)\dots(Y-1)$ contains $MBR_{(k,l)}$ whose centroid is north, west or northwest of the centroid for $MBR_{(i,j)}$.

A southwest test is not required because it is the inverse of the northeast test.

Figure 9 depicts the node validity for the set of MBRs in Figure 9(a). The node in Figure 9(b) is valid for the set of MBRs. MBR 1 is in location (0,0). For MBR 2 to be in location (1,0), its centroid must be, and is, southwest of the centroid for MBR 1. Similarly, for MBR 3 to be in location (0,1), its centroid must be located northwest of the centroid for MBR 1. For MBR 4 to be in location (1,2), its centroid must be northeast of both MBRs 1 and 3, and northwest of MBR 4. All conditions are satisfied, so the node is valid. The node in Figure 9(c) is invalid. In this case, MBR 4 is in location (1,1), which means that its centroid is southeast of the centroid for MBR 3. However, MBR 4 is northeast of MBR 3. Therefore, the node is invalid.

Figure 10 shows a 2DR-tree that preserves all spatial relationships for the given objects (taken from [4]). In the leaf node (m4,m3), the centroid of m3 is located southeast from the centroid of m4, so m3 is located east of m4 in the

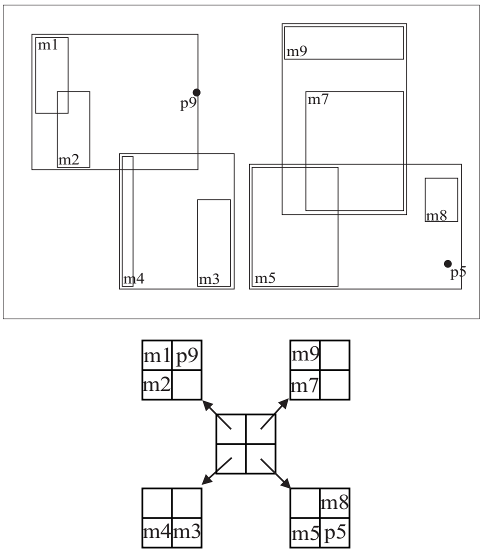

Figure 10: Order 2*2 2DR-tree

node. In node (m5,p5,p8), p5 is located southeast of the centroid for m5, and the centroid for m8 is located northeast of the centroid for m5 and northwest of p5. Therefore, p5 is stored east of m5 while m8 is stored northeast of m5 and north of p5. The remaining two leaf nodes, (m7,m9) and (m2,m1,p9) preserve the spatial relationships between their objects. Spatial relationships between subregions in the root are also preserved.

5 Spatial Indexing in TIP

This section addresses the first of the four issues identified in Section 3. We first present the proposed TIP indexing structure, followed by our preliminary approach for processing an incremental range query using the index. We also present extensions to the preliminary approach, and other issue under consideration.

5.1 The TIP Index

The overall index for TIP will consist of three parts. The first is the spatial index. A 2DR-tree [8, 9, 10] will be constructed using the coordinates from the Gazetteer. The second part will consist of locations in the Gazetteer. Each point in a leaf node will reference a list of all locations at the point. For example, one point can have Hamilton Gardens, Hamilton, and Waikato associated with it. Also, for each location, the information provided in the Gazetteer, such as population, can also be stored with the location. The third part consists of a

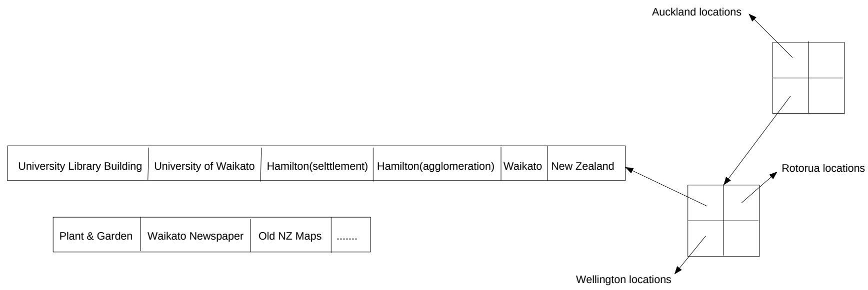

Figure 11: Spatial indexing for TIP

list of Greenstone collections. Each location will reference a corresponding list of collections.

Figure 11 depicts a portion of a TIP indexing structure that includes Hamilton and the area surrounding it. Focusing on the leaf level of the 2DR-tree, one node contains 3 pairs of `<x,y>` coordinates. One pair references a list of all locations accessible from the point. These locations are the spatial objects and concepts depicted back in Figure 7. Each location will reference a list of Greenstone collections that are related to the location. For example, 'Hamilton' will refer to the 'Plant and Garden' and the 'Old NZ maps' collections.

5.2 Incremental Search

Our preliminary approach to the incremental range query extends the existing 2DR-tree binary search strategy (see [8, 9] for details) with node traversal. The incremental range query works as follows. When the user begins their trip, the coordinates of their initial location are provided to the TIP index. A binary search is performed to locate both a node and a location within the node. This serves as the starting point for the incremental the range query. From here, the query is guided by the user's movements. For example, if the user moves in the northeast direction, the range query will proceed in the northeast direction within the node. Likewise, if the user moves in a southeast direction, the range query will proceed to the east within the node. Similarly, the user's movements will guide the movement of the range query within the node.

For example, we consider the user at the University of Waikato. When they switch to the TIP/Greenstone bridge, the coordinates of their GPS location are provided to the spatial index. A binary search is performed, which leads to the list of locations given in Figure 11. After the user selects Hamilton, the corresponding lists of Greenstone collections are retrieved and displayed to the user. If the user proceeds southeast towards Rotorua, then the range query will move to the east within the original leaf node. Then, all locations accessible from the next coordinates can be retrieved and displayed to the user. Similarly, the set of locations related to Wellington can be retrieved simply by traversing the leaf node to the south from the coordinates for Hamilton.

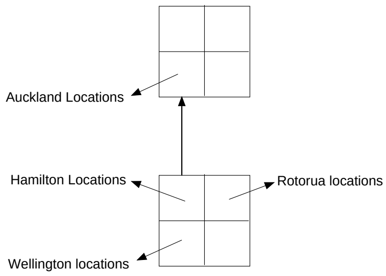

Figure 12: Link between two leaf nodes

This guided search is possible because the points in the 2DR-tree are organized based on their relationships with each other. In addition, as long as the incremental range remains within the node, we do not need to perform a new query from the root of the tree every time the user moves.

5.3 Extensions

Incremental search is expected to work well if the user stays in a small area. Unfortunately, if the user travels begins to travel any significant distance, additional binary searches will be necessary because the bounds of the current node will be reached. For example, if our user travels east or southeast of Rotorua, or northwest of Hamilton towards Auckland, a new binary search would need to be performed in order to locate a new starting point for the incremental range query.

One extension to the spatial index for TIP that is under consideration is to relate the leaf nodes with links. This idea is similar to the leaf node links used in the $B^+ - tree$ [3]. Figure 12 depicts an example link between two leaf nodes. Here, if the user proceeds northwest from Hamilton to Auckland, the range query can be moved to the next node by following the link. This would eliminate the need for an entire binary search to be performed.

5.4 Other Considerations

We finish this section with other issues that must be addressed in the implementation of the strategy outlined above.

5.4.1 Leaf Node Links

When adding links between leaf nodes, it must be done so that spatial relationships are preserved between points that exist across different nodes. Because the TIP index incorporates the 2DR-tree, all spatial relationships are preserved between objects or points in each node at the leaf level, and between all regions that contain subjects of objects or points in each node at the non-leaf levels. Therefore, the addition of links at the leaf level looks promising.

However, the current method of constructing a 2DR-tree is by inserting each object or point one at a time. Although spatial relationships between objects within a node are maintained, the spatial relationship between two objects that belong to two different nodes cannot be easily determined. We suggest a different technique for constructing the TIP index, namely a bottom-up tree construction strategy. This strategy will require a strategy for clustering coordinates that considers the spatial relationships between them. Then, the clustered coordinates are organized into a 2DR-tree. It is expected that the resulting tree structure will lend itself to the linking of its leaf nodes.

5.4.2 Size of a Range Query

The proposed incremental range query for the TIP index assumes that the range query does not extend past the bound of the current leaf node that is being traversed. This can be guaranteed for a point query (i.e. a range query of zero area). However, as mentioned earlier, the user rarely, if ever, stands on top of the area that they needs information about. Therefore, range queries will be used in the TIP index. The problem that needs to be addressed is how to ensure that information across multiple leaf nodes is retrieved. Initially, this will be accomplished with the initial binary search. The problem arises when traversing across a node - it is possible that the edge of the range query travels outside the edge of the node. We expect to address this after leaf node links have been added to the index.

5.4.3 Node Traversal

If a user is traveling north, south, east, west, northeast or southwest, the corresponding movement in the current leaf node is straightforward. If the user is travelling northwest or southeast, the next move is not as obvious, since for each there are three options to choose from. For example, if the current location in a node is $(1, 1)$, and the user wants to move southeast, a move to $(1, 0)$, $(2, 0)$, or $(2, 1)$ are all possible. Therefore, ways to determine which direction to traverse must be considered.

6 Conclusion

In this paper, we analysed the co-operating systems TIP and Greenstone for challenges for efficient access to context-aware information. We gave specific focus to location-based access to spatial data. Four challenges were identified: (1) 2D-range and incremental-range queries, (2) context dimensions, (3) distributed indexes, (4) mobile devices. We presented our 2DR-indexing extension and approach for incremental range queries to address the first challenge. The challenge of further context dimensions will be included in the near future by extending the 2DR-tree into a nDR-tree.

References

[1] R. Bayer and E. McCreight. Organization and maintenance of large ordered indices. *Acta Informatica*, 1:173–188, 1974.

[2] J. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18:509–517, 1975.

[3] D. Comer. The ubiquitous B-tree. *ACM Computing Surveys*, 11:121–37, 1979.

[4] V. Gaede and O. Günther. Multidimensional access methods. *ACM Computing Surveys*, 30:170–231, 1998.

[5] X. Gao. Audio icon and audio books service in tip. Master's thesis, 2007. Department of Computer Science, University of Waikato, Hamilton, New Zealand.

[6] A. Hinze and G. Buchanan. The challenge of creating cooperating mobile services: Experiences and lessons learned. In *Proceedings of the Twenty-Ninth Australasian Computer Science Conference (ACSC 2006)*, pages 207–215, 2006.

[7] A. Hinze, X. Gao, and D. Bainbridge. The tip/greenstone bridge: A service for mobile location-based access to digital libraries. In *Proceedings of the European Conference of DigitalLibraries (ECDL)*, pages 99–110, 2006.

[8] W. Osborn. *The 2DR-tree: a 2-dimensional spatial access method.* PhD thesis, University of Calgary, University of Calgary, 2005.

[9] W. Osborn and K. Barker. Searching through spatial relationships using the 2DR-tree. In *Proceedings of the 11th International Conference on Internet and Multimedia Systems and Applications*, 2006.

[10] W. Osborn and K. Barker. An insertion strategy for a two-dimensional spatial access method. In *Proceedings of the 9th International Conference on Enterprise Information Systems*, 2007.

[11] P. Rigaux, M. Scholl, and A. Voisard. *Spatial databases: with application to GIS.* Morgan-Kauffman, 2001.

[12] S. Shekhar and S. Chawla. *Spatial databases: a tour.* Prentice Hall, 2003.

[13] I. H. Witten and D. Bainbridge. *How to Build a Digital Library.* Morgan Kaufmann, 2002.