LNCS 2750, pp1-9, October 2003.
(URL: http://www.springerlink.com/link.asp?id=22qlbrtavtyu1vk9)

Data Modeling for Mobile Services in the Real World

Christian S. Jensen, Torben Bach Pedersen, Laurynas Speičys, and Igor Timko

Department of Computer Science, Aalborg University, Denmark

Abstract. Research contributions on data modeling, data structures, query processing, and indexing for mobile services may have an impact in the longer term, but each contribution typically offers an isolated solution to one small part of the practical problem of delivering mobile services in the real world. In contrast, this paper describes holistic concepts and techniques for mobile data modeling that are readily applicable in practice. Focus is on services to be delivered to mobile users, such as route guidance, point-of-interest search, road pricing, parking payment, traffic monitoring, etc. While geo-referencing of content is important, it is even more important to relate content to the transportation infrastructure. In addition, several sophisticated, integrated representations of the infrastructure are needed.

1 Introduction

The integration of location-based content and the reuse of content across multiple services will be central to the cost-effective delivery of competitive mobile services as well as to the rapid development of new mobile services. Appropriate data modeling is therefore essential to the delivery of quality mobile services.

Knowledge of the user's location is central to many mobile services. Knowledge of the transportation infrastructure in which the user is moving is often also essential to a mobile service. A general-purpose foundation for the delivery of location-based services requires that multiple, integrated representations of the infrastructure are available. In addition, content (also termed business data) must be geo-referenced and must be positioned with respect to the infrastructure.

The paper offers the reader insight into the real-world challenges to data modeling for mobile services, and it describes an approach to data modeling that meets the challenges. The approach is being used by the Danish road directorate and Danish company Euman [2]. It is our hope that this paper will shed light on the application domain of mobile services, will demonstrate some of the complexity of the inherent data management problem, and will inform future research.

Much related research in the area of computer science makes simple assumptions about the problem setting. The transportation infrastructure is often not taken into account. As a result, the notions of proximity used are inappropriate

T. Hadzilacos et al. (Eds.): SSTD 2003, LNCS 2750, pp. 1–9, 2003.

to many services. Further, only a limited range of services can be supported. Other research reduces a transportation infrastructure to a graph, but does not consider other aspects of the infrastructure. The present paper takes steps towards integrating these contributions, thus taking first steps towards making the advances from the scientific domain relevant in practice. In the industrial domain, linear referencing [4, 6] has been used quite widely for the capture of content located along linear elements (e.g., routes) in transportation infrastructures. In addition, a generic data model, or ER diagram, has been recommended for the capture of different aspects of entire transportation infrastructures and related content [5], and a number of variants of this model have been reported in the literature (see [3] for further references). The data model presented in this paper employs linear referencing, and it improves on other data models in several respects. Most notably, it is the only model known to us that integrates different representations of a transportation infrastructure via a purely internal model of the infrastructure. In a sense, this approach is a "geographical generalization" of the use of surrogate keys for the management of business data.

The paper is structured as follows. The case study in the next section describes key data modeling challenges. Sect. 3 describes the data modeling approach that meets the challenges. Finally, Sect. 4 summarizes the paper.

2 Case Study

We proceed to illustrate data modeling challenges posed by mobile services.

2.1 Requirements

The kinds of queries to be supported induce requirements. A simple type of query computes the distance from a user's current position to a point of interest, such as an art gallery or a nature park. Another type of query concerns route planning, where, e.g., the user wants to retrieve the "shortest" route to a certain point of interest while passing one or more points of interest enroute. Yet another type of query retrieves the "nearest" point of interest such as a particular type of store or gas station. The terms "shortest" and "nearest" may be given several meanings. For example, they may be based on Euclidean distance or the anticipated travel time along the road network.

Queries such as these depend on two types of data: the transportation infrastructure and the "remaining" data. The remaining data is sometimes termed business data or content. Examples include the points of interest mentioned above.

To position a service user wrt. the infrastructure, it is necessary to geo-reference the infrastructure. To perform route planning, a graph-like representation of the transportation infrastructure is helpful. To position some content, e.g., information about accidents, wrt. the infrastructure, a so-called known-maker representation of the infrastructure is required. Next, in order to support queries such as those listed above, the business data must be geo-referenced as well as positioned wrt. all the different representations of the infrastructure. The implication is that the infrastructure representations must be interrelated, meaning that it must be possible to translate from one representation to another.

2.2 Content

Content generally falls into one of two categories. First, *point data* concerns entities that are located at a specific geographic location and have no relevant spatial extent. This type of data is attached to specific points in the transportation infrastructure. Second, *interval data* concern data that are considered to relate to a *part* of a road and are thus attached to intervals of given roads. Interval data can be categorized as being (1) *overlapping versus non-overlapping* and (2) *covering versus non-covering.*

Non-overlapping, covering content includes speed limits and road surface type. Non-overlapping, non-covering content includes u-turn restrictions and road constructions, as well as more temporary phenomena such as traffic congestion and jams. Examples of overlapping, non-covering content includes tourist sights. A scenic mountain top and a castle may be visible from overlapping stretches of the same road. Other part of roads have no sights. Another example is warning signs. Overlapping, covering content include service availabilities, e.g., a car repair service may be available from some service provider anywhere, and several repair service providers may be available in areas.

2.3 Transportation Infrastructure

We proceed to describe the actual infrastructure that is to be modeled. To be specific, we consider two consecutive "intersections" along a single road. Aerial photos are given in Fig. 1. The road we consider first stretches from West to East (a), then bends and goes from South-West to North-East (b). We describe each intersection in turn.

While our road is generally a bidirectional highway with one lane in each direction and no median strip dividing the traffic, the first intersection introduces a median and includes two bidirectional exit and entry roads. Major concerns underlying this design are those of safety and ease of traffic flow. Vehicles traveling East, i.e., using the lower lane of the highway, must use the first road to the right for exiting. A short deceleration lane is available. The second road connected to the right side of the highway is used by vehicles originating from the crossing road that wish to travel East on the highway. A short acceleration lane is available. A similar arrangement applies to the highway's upper lane.

At the second intersection, in Fig. 1(b), the highway has two lanes in each direction, a median, and four unidirectional exit/entry lanes. A vehicle traveling North-East, i.e., using the right lane of the highway, can decelerate in the long exit lane, while North-East bound vehicles must enter and can accelerate via the long right entry lane.

It should be clear that a transportation infrastructure is not just a mathematical graph. While some aspects may be described as a directed, non-planar

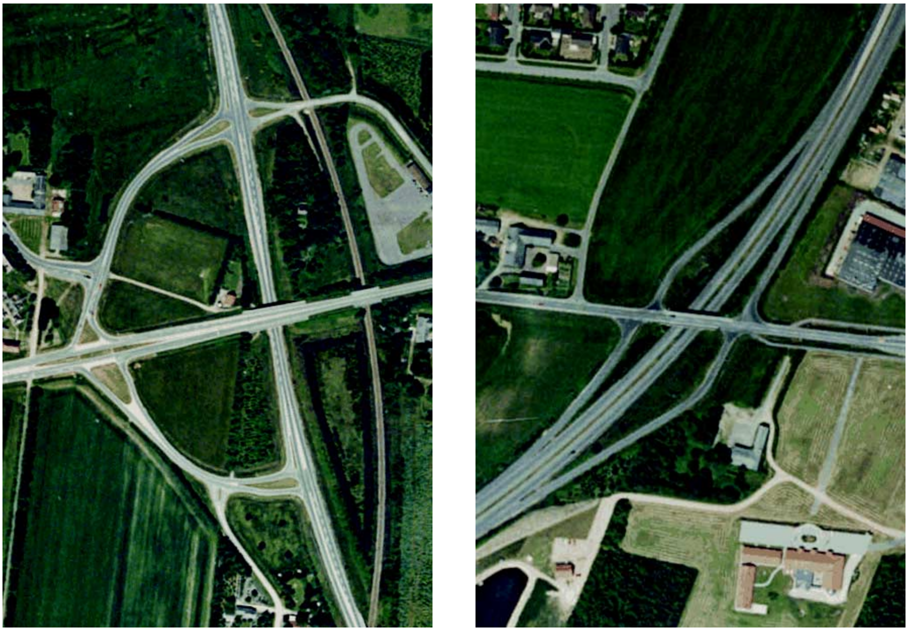

(a) First Intersection (b) Second Intersection

Fig. 1. Aerial Photos of Road Intersections

graph, many other aspects are left unaccounted for by such a simple representation.

3 Data Modeling

The data model provides three external, user-accessible, representations of the transportation infrastructure, namely the *kilometer post*, *link-node*, and *geographic* representations, which are connected by an internal *segment* representation. The core of the data model is given in Fig. 2. When referring to this figure, we use capital letters and italics, e.g., ROAD and *r_id*, to denote table names and attributes, respectively.

3.1 Kilometer-Post Representation

The kilometer-post representation (the most commonly used type of known-marker representation) is used for road administration. It is convenient for relating a physical location to a location stored in a database and vice versa. Location is expressed in terms of the road, the distance marker on the road

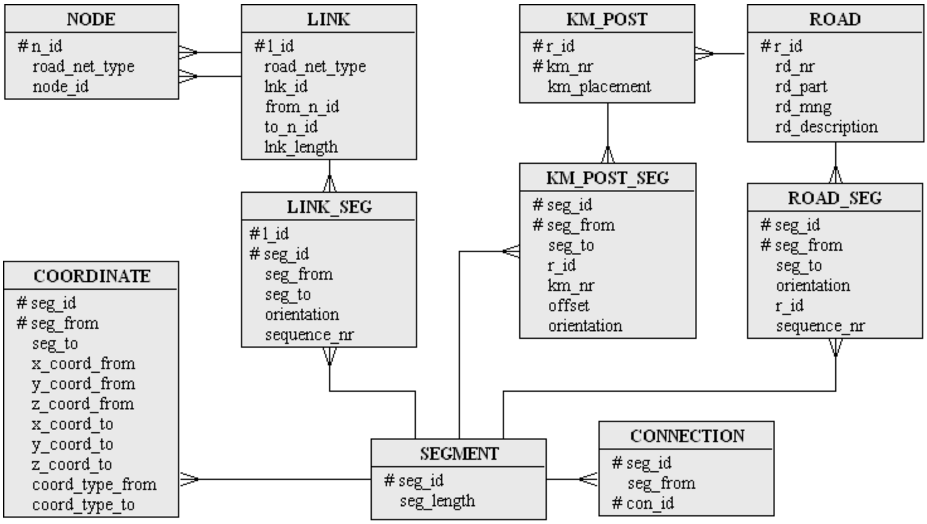

Fig. 2. Core Data Model

(e.g., kilometer post), and the offset from the distance marker. The representation is used by administrative authorities for collecting and utilizing data on field conditions, i.e., for entering content into the system. Primitive technological means, such as a simple measuring device and map and a ruler, suffice for identifying a position on a road, rendering the use of the representation cost effective and thus practical.

Table ROAD captures the administrative identification of roads, using the following attributes: the internal surrogate road id *r_id*, the external road number *rd_nr*, the road part *rd_part* (physically separate lanes, etc.), the road management authority *rd_mng*, and the external name of the road *rd_description*. Table KM_POST captures information on road distance markers, namely the full kilometer distance *km_nr* from the start of road *r_id*, possibly with an offset *km_placement* from the physical marker (used, e.g., when the logical marker is in the middle of an intersection). The kilometer-post representation of the case study is seen in Fig. 3. Imaginary markers (with residual meter offsets from the full kilometer in parentheses) are used to mark beginnings of road parts.

3.2 Link-Node Representation

The link-node representation is based on the concepts of weighted undirected and directed mathematical graphs. A node is a road location with a significant change of traffic properties, e.g., an intersection. A link is a route that connects two nodes. Such a representation abstracts away geographical detail, but preserves the topology of the transportation infrastructure. For this reason, node-link representations are suitable for tasks such as traffic and route planning. The former task refers to (re)designing road networks taking traffic flows

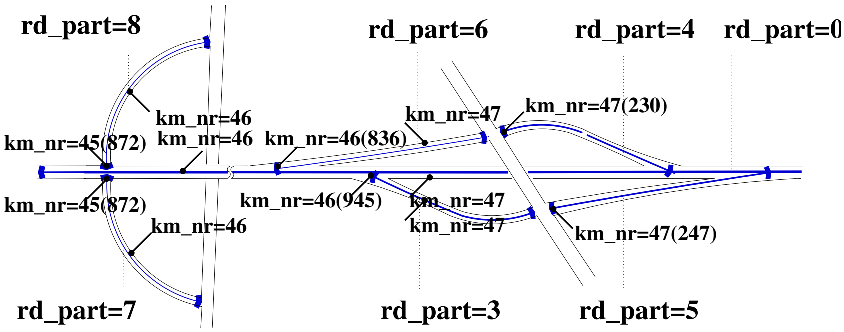

Fig. 3. Kilometer Post Infrastructure Representation

into consideration. In this case, undirected graphs are sufficient. The latter task refers to finding traversable routes in road networks that satisfy certain criteria. Directed graphs that capture traffic directions are appropriate for this task. The representation allows several *road network types*, e.g., a coarse-grained one for traffic planning and a fine-grained one for route planning.

Table NODE captures nodes, i.e., the internal surrogate id *n_id*, the road network type *road_net_type*, and a generic *node_id* that spans all road network types. Table LINK captures a link from start node *from_n_id* to end node *to_n_id* of length *lnk_length*. Links also have identifier attributes analogous to those of nodes, i.e., the *road_net_type*, the generic *lnk_id*, and the unique *l_id*.

3.3 Geographical Representation

The geographical representation captures the geographical coordinates of the transportation infrastructure. The coordinate representation enables users to directly reference location rather than measure distances along roads from certain locations, such as kilometer posts. Additionally, the representation is used by geography-related applications, such as cartographic systems or certain GISs, that operate on coordinate data.

Table COORDINATE captures a three-dimensional point given by (*x_coord_from*, *y_coord_from*, *z_coord_from*) on the center-line of a road. Several levels of detail are possible; these are captured by attribute *coord_type_from*. Attributes *x_coord_to*, *y_coord_to*, *z_coord_to*, and *coord_type_to* indicate the next 3D point on the segment. They are redundant and used to enhance the efficiency of query processing. Attributes *seg_id*, *seg_from*, and *seg_to*, map the two points of the tuple to the segment representation, described next.

3.4 Segment Representation

The internal segment representation models an infrastructure as a collection of segments that intersect at connections (locations where there is an exchange of

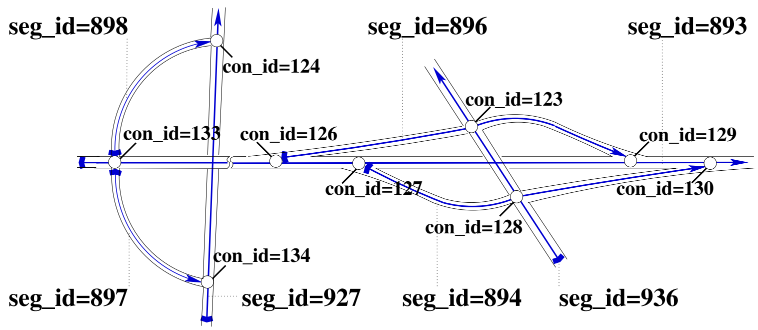

Fig. 4. Segment-Based Infrastructure Representation

traffic). This representation preserves the network topology and captures the complete set of roadways.

The positions of any content (e.g., speed limits, accident reports) are given by references to segments. In addition, the external representations of the infrastructure are mapped to the segment representation in a way that establishes one-to-one mappings between all the representations, i.e., the segment representation is the *integrator* of the different representations. For this reason, the segment representation is used by content integrators for efficient content position maintenance and translation between the external representations. The segment representation is purely internal to the system.

Table SEGMENT captures segments with unique id *seg_id* and length *seg_length*. A row in table CONNECTION describes that segment *seg_id* intersects with a connection point identified by *con_id*. The intersection of the segment with the connection point occurs at *seg_from* units from the start of the segment. The segment representation of the case study is illustrated in Fig. 4.

3.5 Interrelating the Representations

As pointed out earlier, we need the ability to translate among the different external infrastructure representations. This is achieved by connecting the kilometer-post representation and the node-link representation to the segment representation, which is already integrated with the geographical representation.

In our model, these connections are recorded by the tables KM_POST_SEG and ROAD_SEG for the kilometer-post representation and by table LINK_SEG for the link-node representation. We consider the latter two tables only.

As for the kilometer-post representation, a row in table ROAD_SEG relates (a section of) a road part *r_id* to a part of a segment *seg_id*. The (section of the) road part positioned by a row corresponds to the section of the related segment with end points at *seg_from* and *seg_to* units from the start of the segment. The attribute *orientation* indicates whether or not the directions of linear referencing along the segment and along the road part coincide. Attributes *rd_id* and *seg_id* are foreign keys that reference the tables ROAD and SEGMENT, respectively.

Further, since road parts do not overlap, the pair (*seg_id*, *seg_from*) is the primary key of the table. Finally, *sequence_nr* is an ordinal number of the road part section. It is used for distinguishing among different sections of the same road part and to "reconstruct" road parts from a collection of segment sections.

Note that table KM_POST_SEG alone fully defines the relation between the kilometer-post and other representations. Table ROAD_SEG is merely included for query efficiency—it contains redundant information.

As regards the link-node representation, a record in table LINK_SEG positions (a section of) a link *l_id* within a segment *seg_id*. The attribute *orientation* indicates whether the directions of linear referencing along the segment and of the link coincide. The other attributes have the same semantics as do their counterparts in table ROAD_SEG. Attributes *l_id* and *seg_id* are the foreign keys that point to the primary keys in tables LINK and SEGMENT, respectively.

3.6 Integration of Content

In our model, each type of content is allocated a separate descriptive table. For example, tables ACCIDENT and SERVICE describe instances of road accidents and car repair service providers, respectively. Further, each type of content is associated with a table that positions the content with respect to the infrastructure in the segment representation, e.g, a table ACCIDENT_SEG captures accident position and SERVICE_SEG captures service availability ranges.

The principles of positioning interval data with respect to segments are described above. In particular, the tables for interval content have schemas similar to those of tables ROAD_SEG and LINK_SEG. The positioning of point data is analogous to the positioning of connections (see above).

The same content must be accessible via different infrastructure representations. Given a type of content, a (possibly materialized) view can be created for each necessary external representation to allow easy and fast access.

4 Summary

Mobile, location-based information services, such as traffic, weather, tourist, and safety-related services, are seen as a very important new type of application. Such services depend on the availability of geo-related content, on multiple representations of the geographical infrastructure in which the users are traveling, and on the integration of the content with the infrastructure. This paper presents a case study with requirements, and it covers data modeling in response to the requirements, concentrating on infrastructure representations: the kilometer post, node-link, geographical, and segment representations, and on the integration of these with each other and the related content.

We hope this paper will expose to the data management research community some of the challenges faced when building a system capable of delivering high-quality mobile services in practice.

Danish company Euman is building a content integration and delivery system that uses an approach to data modeling that is quite similar to the one described in this paper.

Acknowledgments

We thank Christian Hage at Euman A/S for many discussions on the paper's topics. This work was supported in part by two grants from the Danish Centre for IT Research. The photos in Fig. 1 were made available by COWI, who hold the rights to the photos.

References

[1] P. Djernæs, O. Knudsen, E. Sørensen, and S. Schrøder. VIS-brugerhåndbog: Vejledning i opmåling og inddatering. 58 pages, 2002. (In Danish)
[2] Euman. `http://www.euman.com` 1
[3] C. S. Jensen, T. B. Pedersen, L. Speicys, and I. Timko. Integrated Data Management for Mobile Services in the Real World. *VLDB Conf.*, 2003 (to appear) 2
[4] C. Murray. Oracle Spatial User Guide and Reference, Release 9.2. *Oracle Corporation*, 486 pp., 2002 2
[5] NCHRP. A Generic Data Model for Linear Referencing Systems. *Transportation Research Board, Washington, DC*, 28 pp., 1997 2
[6] P. Scarponcini. Generalized Model for Linear Referencing in Transportation. *GeoInformatica* 6(1):35–55, 2002 2