THE ONTARIO STANDARD LABELLED ROAD NETWORK: INTEGRATION OF EMERGENCY SERVICES AND OTHER USER NEEDS INTO A CONCEPTUAL DATA MODEL[1]

Valerian T Noronha
Digital Geographic Research Corporation
Mississauga, Ontario, Canada
(416) 410-2030
noronha@dgrc.ca

Michael F Goodchild
National Centre for Geographic Information and Analysis
University of California
Santa Barbara, California 93106
(805) 893-8049
good@ncgia.ucsb.edu

Abstract

Probably the most dramatic application for digital street networks is in emergency response. Information quality, especially accuracy and currency, is critical. The cost of compiling and maintaining high quality digital networks is substantial, and cannot be borne by emergency services alone. Other potential users of network data may be willing to share the costs if their information needs are satisfied. This points to the need for a standard data model, that (a) addresses the information content requirements of a variety of user groups, and (b) facilitates data exchange. Our study of Ontario's Emergency Health Services and other user groups identified a range of needs. The standard data model presents a broad view of networks at four information levels, for applications such as production cartography, position reckoning, address matching and route optimization. A broad base of users, from commercial cartographers to fast-food outlets, can expect to benefit from network data that conform to this model.

---

[1]Earlier versions of this paper appeared in the Proceedings of the Fourth Canadian Conference on GIS, Ottawa, March 1992 (Energy, Mines and Resources Canada); and the Proceedings of the GIS Seminar, Toronto, November 1991 (Ministry of Natural Resources, Ontario).

1. INTRODUCTION

One of the most compelling justifications of a digital information system would be the need to access and process crucial data under pressure of an emergency. The need for *geographic* information processing in these situations is almost universal. The physical delivery of assistance requires information on the precise location of the emergency, and the most effective route to the scene.

The Emergency Health Services (EHS) branch of the Ontario Ministry of Health is responsible for delivering ambulance service to all of Ontario except Metro Toronto. Unlike police and fire, which are administered by local governments, EHS must deal with the entire province, and deliver service equitably both in the Golden Horseshoe megalopolis and in the remote north.

In 1990 EHS began to explore the creation of a digital road network data base suited to emergency dispatching, tentatively termed the Ontario Standard Labelled Road Network (OSLRN). It would be an authoritative source of current road data for the province. Given that a road in this discussion is any highway or street designed for a motor vehicle, the magnitude of this task is clearly enormous. EHS assembled experts from several public and private bodies, to study the range of potential applications of network data, and to examine the possibilities of data exchange and cost sharing. The group included members such as the Milk Marketing Board and the Ministry of Education, who schedule vehicles (milk trucks and school buses respectively) along provincial roads, and have an interest in computer assisted routing and scheduling.

It requires a high order of political cooperation and technical coordination to realize the OSLRN data base. A first, modest objective would be a detailed definition of a conceptual road network model, that would address the needs of most of the anticipated users of the data. If the model was sufficiently generic that it could be adopted by municipalities, system vendors and data custodians, then the technical barriers to data exchange would be largely eliminated. The group commissioned the authors to design the OSLRN data model. The formal model definition is contained in other documents; this paper approaches the matter from an applications viewpoint.

2. APPLICATIONS OF ROAD NETWORKS

The SLRN design process began by identifying the principal potential users of medium scale (1:50,000 to 1:10,000) centreline network data, and the information sets required by each. A centreline (as opposed to kerbline) data base represents a street by a single polyline running approximately through the centre of the street. Kerbline files are useful only for large scale (larger than 1:10,000) municipal type applications.

Present and potential users of centreline files include cartographic firms, municipalities (scheduling of garbage collection, street cleaning and snow ploughing), utility firms, wholesale distributors (soft drink companies, supermarket chains), political parties, market researchers, direct mail companies and the post office, all home delivery (pizza, newspaper, department stores), and emergency response units. There are four basic application areas:

Cartography (base maps) Medium scale topographic sheets usually show roads only in terms of positional alignment, with the class of road sometimes indicated by line symbology. Primary highways are usually annotated, whereas urban streets are not named. This level of information serves only to indicate the existence of a network.

Basic reckoning and navigation (street maps) From the viewpoint of everyday applications, a more useful network representation is the one found in mass-market street maps. By simple visual correlation with landmarks, a disoriented user can gather enough information to determine the general location of a street, e.g. Yonge Street, and can trace a suitable route to a destination.

Address matching Basic navigation as described above is sufficient for occasional users of address data. Professional users, such as emergency crews, delivery personnel and market researchers, must usually be able to geocode an address, e.g. 127 Yonge Street, to block face resolution. In other words, they require coordinates for the address, that are accurate within no more than a few hundred metres. In applications such as school busing, it is important to know on which side of the street a given address falls. Specialized map books and city guides may print civic numbers at major intersections; detailed address ranges and coordinates for block faces are available in digital form for major Canadian cities, from Statistics Canada's Area Master Files (AMFs).

Routing and scheduling In a digital environment, network data can be used to determine optimal paths, and to schedule flows with capacity constraints. This application requires topological networks, in which the relationship of each road to each of its intersecting neighbours is explicitly specified, and impedance measures (e.g. distance) are attached to each road segment. Rural routing is usually straightforward, since congestion and traffic control are not significant issues. On the other hand an urban routing system must ideally consider stop signs, traffic light programming and instantaneous congestion measures, in determining the impedance of a route. The cost of providing current data on this scale is prohibitive, hence one must be realistic in one's expectations of the scope of automated routing.

A network model should address each of the above applications, and should also be designed to facilitate error detection and correction, and update.

3. A BRIEF REVIEW OF DATA MODELS

Digital street data have been in circulation for several years now, notably through the AMF in Canada, and the US Census's DIME and TIGER files. These data bases were originally designed to assist census enumerators, and the data models may not be well suited to popular modern applications such as those detailed above.

DIME/TIGER The DIME files were designed to store urban polygons (city blocks) and streets simultaneously in a unified data structure. In the attempt to maintain topological integrity of the polygons, DIME viewed streets essentially as polygon boundaries. When two streets crossed paths, there was no distinction between traffic intersections and non-grade crossings such as overpasses. TIGER, designed for the 1990 US Census, is a major reconstruction of the original DIME structure. Yet in both cases the primary entity is a road segment between two intersections. A major road, such as Yonge Street in Toronto, is therefore fragmented into numerous short links. In general a fragmented model contains data redundancies and is inefficient for both storage and access — although performance ultimately depends also on the data structure and implementation.

AMF The Area Master File treats a street as a single continuous entity, listing intersecting streets as they occur along its path. Attributes such as address ranges are logically attached to each segment.

The crucial problem common to AMF and DIME/TIGER is that the building block is the urban block face (Statistics Canada has recently improved the AMF structure; it is now based on topological chains). Attributes associated with the network, such as pavement condition, the number of lanes, or the location of point landmarks, must therefore be coded with the block face (or topological chain) as the basic unit of resolution. This is clearly inadequate, and especially where block faces are long, this representation can result in serious analytical errors. Alternately, the resolution must be improved, and each street must be broken wherever *any* attribute changes. The result is a file that is disproportionately large, given its information content.

4. THE SLRN MODEL

The SLRN model is reactive rather than proactive, in that it is a response to the way networks are conceived and used, rather than a prescription for instituting radically different methods. The model is built on the concept of ‘‘dynamic segmentation,’’ that has recently been discussed in the GIS literature. A person normally has a mental construct of a road as a logical and coherent entity — this is equally true of neighbourhood crescents as it is of cultural landmarks such as Yonge Street. The SLRN maintains the integrity of the road, and attaches auxiliary information such as address ranges and intersections when appropriate.

The model is designed to address all the four principal application areas of street networks identified above, namely, cartography, position reckoning, address matching and routing. Accordingly, there are four information levels in the model.

Level 1

The term **thoroughfare** is used to describe the general class of motorable surfaces in a network. The basic SLRN entity is a **road**, a continuous length of thoroughfare which is distinguished from the rest of the network by virtue of having a defined beginning and end. The model is flexible as to how the beginning and end of the road are defined. Consider for example Toronto’s Yonge Street. It runs from Queens Quay in southern Toronto to Steeles Avenue, the northern limit of the city (Figure 1). A data base for Metro Toronto may choose to define Yonge Street as a single entity from Queens Quay to Steeles, or may break it into say two sections, the first from Queens Quay to Eglinton, the second from Eglinton to Steeles. The partitioning would usually be designed to correspond with existing spatial divisions or ‘‘tiles’’ in the data base, such

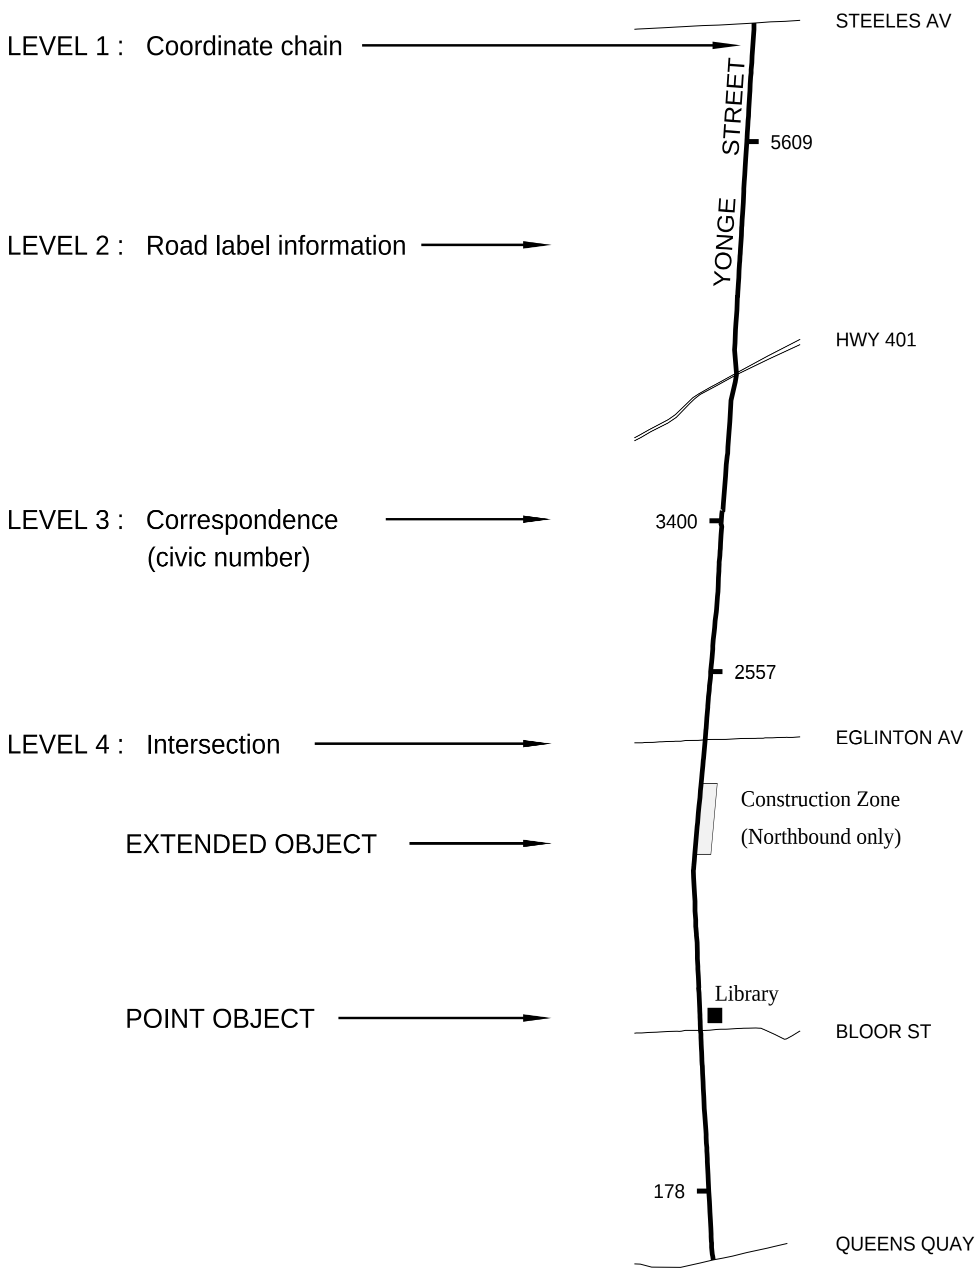

Figure 1: SLRN representation of Yonge Street, Toronto

as map sheet or civic ward boundaries. In an extreme case the road may be fragmented into several hundred block faces, but that would clearly negate the advantages of the SLRN model.

Points along the road are addressed in terms of their **offsets**, i.e. their distance from the beginning point. The name of the road is not stored at Level 1, yet the fact that it is continuous, and connects a defined pair of end points, does lend the road an implicit identity.

From an applications viewpoint, Level 1 data are in effect coordinate chains for unlabelled cartographic plotting such as in topographic maps. Since the road is a continuous entity, symbology and line weight are constant for the entire stretch of road.

Level 2

The second information level is the labelled road. In reality a road may have an alias, or a physically continuous road may have more than one name, either along its entire course or in different stretches. There may be more than one road with the same name (e.g. King Street), usually assigned by different municipalities.

One or more road names may be attached to any stretch of a SLRN road, each stretch being defined by a start-offset and end-offset. Each name record has fields for the naming authority and the civic numbering system. A Lot-and-Concession nomenclature may therefore coexist with a civic addressing system on the same physical road.

Level 2 information is sufficient for automating production of the network layer (including annotation associated with the network) in street maps. In a production environment one may need to store road class information (e.g. primary or secondary highway; major or minor urban street). The model does not explicitly provide for a road class field, since the assignment of classes is ultimately subjective. If the road is formally assigned a class by the naming authority (e.g. The King's Highways in Ontario), this applies to the entire section of road covered by that name, and could be indicated in a Level 2 name field without violating the fundamental character of the model. On the other hand, if class is a judgement of the importance of the road, based on the number of lanes, traffic capacity or average speed, this information would often vary along the course of the road, and is better stored as an Extended Object (see below).

Level 3

Civic address data are usually made available as block face address *ranges*. It is a simple linear interpolation problem to estimate the coordinates of say 127 Yonge Street, given coordinates for 100 Yonge Street and 200 Yonge Street. With long block faces, however, the accuracy of interpolation relies on the questionable assumption that addresses are assigned in regular linear fashion. Accuracy would be improved if coordinates were known for some intermediate addresses in a block; the traditional model does not allow for these to be stored.

The SLRN introduces the concept of a **correspondence**. This is simply a logical mapping of a random offset to its known civic number. By departing from the block face as a unit of resolution, this structure permits correspondences anywhere along the road. Since correspondences may easily be established at beginnings and ends of blocks, the SLRN model accommodates conventional address range data.

Note that interpolation between correspondences, say for 127 Yonge Street, yields the *offset* of 127 Yonge. The coordinates of that address are found by subsequent interpolation along the coordinate chain.

Level 4

Although the above levels contain enough information for many useful applications, they store no topological information, and roads may begin and end at arbitrary points. Routing applications require connectivity data at intersections. This can sometimes be inferred from coordinate chains, but since bridges, flyovers and other non-grade intersections complicate the network fabric, the information must be stored explicitly.

In a Level 4 SLRN, each time an intersection is encountered over the course of a road (the **primary** road), the offset is recorded, along with the identifier of the intersecting, **secondary** road. Since the secondary road may cross the primary road more than once, the offset along the secondary road is also stored at each intersecting point. Where two roads cross paths at non-grade intersections, i.e. neither is physically accessible from the other, an intersection is not recorded.

This is sufficient information to construct connectivity tables for simple routing applications. Moreover, with this data model, the procedures that generate connectivity tables can easily be adapted to extract a generalized network (i.e. selected routes, such as major arteries) for macro-level analysis.

The impedance associated with each network link may in the simplest case be calculated as a function of the offset. This may not be appropriate in hilly landscapes, where the offset distance (which is measured along the digitized coordinate chain) may differ substantially from odometer distance. In such cases the impedance, whether odometer distance, average driving time or any other measure, may be recorded as an attribute field (see Associated Objects, below).

Associated Objects

The SLRN model does not recognize any features in the landscape (e.g. gas stations, time zone boundary lines, parks), except to record their interaction with the network. In other words, the SLRN is not a generic spatial data model, and is not designed to store random points, polylines or polygons. There are definable relationships between these geometric features and a network: any point in space can be associated with the *point* on the network that is closest to it; a polyline intersects the network at one or more *points*, or rarely, a line; and a polygon intersects a network over a length of a network *line*.

Accordingly, the SLRN defines **point objects** and **extended objects**. A point object (e.g. gas station, time zone boundary) is an entity associated with a *point* on the network, identified by the road ID and the offset of the point along that road. An extended object (e.g. park) is an entity associated with a *section* of one or more roads, identified by the road IDs, and a start-offset and end-offset along each road. Point objects and extended objects are optional with all levels of SLRN.

Data Revision

Street network data must constantly be revised as new streets and addresses are added to the urban fabric, and as newer and more accurate coordinates are made available. The economic value of the data may be severely compromised if information is obsolete.

The model design explicitly recognizes that coordinates need to be updated periodically, and often selectively. The most likely sources of new coordinates are GPS, photogrammetry, new surveys, and transformations such as NAD27-NAD83 conversion. Quality improvements may take the form of revised or higher-precision coordinates for established points (such as major intersections); better resolution on curves, introducing new points; or additions or deletions of roads or segments of roads.

Additions and deletions are straightforward. For the other two forms of enhancement, update is achieved by first identifying **tie points**, usually all intersections, for which locations are known in both the outgoing and new data bases. Points other than tie points in the outgoing data base will be (a) intermediate nodes in chain definitions, stored as coordinates; and (b) correspondences and associated objects, stored as offsets. Old coordinate chains (a) are discarded and simply replaced by the new versions. Offsets (b) are first converted to coordinates, relative to the tie points in the *outgoing* data base; new offsets are then derived, relative to the new tie points. Offsets will almost always increase with higher resolution data.

5. CONCLUDING REMARKS

The SLRN was designed in early 1991, and is now being circulated among standards agencies in Canada and the United States. It is possible that terminology may be revised, and that other changes may be made in response to public comment.

The apparent shortcomings of the SLRN model are largely intentional. The model deals with the *physical*, not legal aspects of the network. If a road is physically inaccessible from another, as in the case of a simple overpass, the SLRN model does not consider them connected. If the roads intersect at grade, the SLRN blindly records the intersection. There is no provision in the model to indicate whether or not turns are permitted. The reason for this is not technical — Level 4 can easily be broadened to accommodate stop signs and turn tables — but a matter of reality. Legal permissions change daily, and it is rarely that one finds sufficiently current data. At present there is little demand for such information.

Similarly the model can be adapted to store dynamic congestion data. Several cities around the world now have inductive loop sensors embedded in the pavement of freeways and major arteries, to measure traffic volume and velocity. These form the basis of Intelligent Vehicle/Highway Systems (IVHS) of the future. There is speculation that congestion and positional information could be continuously broadcast to on-board vehicle navigation assistance systems, over radio frequencies. This is an expensive proposition, and it will probably be several years before any city implements it on a large scale. Since the SLRN is designed to be reactive rather than revolutionary, detailed model refinement, if required, should be delayed until the direction of these technological developments is clear.

The current SLRN specification is intended for data at an approximate scale of 1:50,000 to 1:10,000, and considers only the centreline of a carriageway. At these scales the representation of complex urban road features (e.g. divided roads, roundabouts) is often a matter of debate. We take the view that these decisions are a matter of scale, cost and benefit. The model readily accommodates either simple or complex (i.e. multiple object) interpretation of these features; it does not attempt to prescribe the most appropriate course of action.

The above examples illustrate the scope and intent of the SLRN model. It is a data *model*, not a data structure. It merely organizes network data into information classes; beyond that it is intended to be a generic and flexible framework, deliberately avoiding specifications on formats, fields and data quality matters. It thus provides a technical foundation for orderly data exchange.

The political dimensions of digital data exchange remain to be resolved. There are undoubtedly costs involved for a data supplier or data user to equip itself to trade digital information. On the other hand, considering that data gathering and input represent the bulk of operating costs in a GIS, the savings on redundant efforts in these areas will easily justify the investment in SLRN compatibility.

The Ontario Ministry of Health and the Ministry of Natural Resources recently embarked on a cooperative pilot venture to pool resources, to exchange data and to develop a Level 3 network for the Ottawa area. Although no software has yet been written around the SLRN model, the SLRN information content standards are temporarily being accommodated within existing systems, albeit at some cost in computing efficiency. The MOH-MNR project involves enhancing MNR’s OBM (Level 1) street data with the information fields specified in Level 3, as required by MOH. Needs such as these drove the development of the model; efforts such as these will ultimately establish its usefulness.