



An Oracle White Paper
February 2011

A Load-On-Demand Approach to Handling Large Networks in the Oracle Spatial Network Data Model

Executive Overview	4
Introduction	5
Oracle Spatial Network Data Model	8
Network Data Model Schema	8
Network Partitioning	10
NDM LOD Architecture and APIs	11
LOD Network Analysis Capabilities	12
Java Representations of LOD Network Elements	13
Software Requirements	14
Using LOD in Network Data Model	14
Creating a Network	14
Partitioning a Network	14
Configuring the Partition Cache	15
Analyzing the Network	17
Modeling and Analysis Enhancements	18
Network Constraints	18
Multiple Cost Support	19
Precomputed Connected Components	20
Dynamic Data Set	21
Points Along A Link	22
Partial Link Paths (Subpaths)	22
Hierarchical Shortest Path Computation	23
Shortest Multiple-Stop Path	24
K Shortest Paths	24
Network Buffer	24
Within-Cost Polygon	25
Minimum Cost Spanning Tree	25
Comparisons Between NDM LOD and In-Memory APIs	25

Approaches	25
Data Model	26
Analysis Functions.....	26
Future Support	26
What Is New in Oracle Spatial 11g Release 2	27
Web-Based LOD Analysis Demo.....	27
Conclusion	29
References.....	29

Executive Overview

Network modeling, management, and analysis are common tasks for enterprise applications in such areas as geographic information systems (GIS), customer relationship management (CRM), social network analysis, and semantic web technologies such as the Resource Description Framework (RDF).

Oracle Spatial 10g introduced the network data model (NDM), which lets users model and analyze networks. In Oracle Spatial 10g, NDM uses an in-memory approach to pre-load the whole network into memory before analysis; however, this approach cannot handle networks that are too large to fit in memory.

To address this scalability issue, we developed a load-on-demand approach (LOD) in Oracle Spatial 11g. Large networks are first divided into manageable parts called network partitions. Only partitions that are needed are automatically loaded during analysis. In this paper we discuss how LOD works and how to use it.

Introduction

The network data model (NDM) helps users analyze network connectivity relationships. It is commonly used in transportation, utilities, life sciences, and semantic technologies. It simplifies network modeling, analysis, and management so that users can focus on application logic. It provides an open, generic data model with many common network analysis capabilities. Application information is separated from connectivity information so that the model can be applied to many network applications without customization. NDM further provides a constraint mechanism, to let users guide analysis based on application rules and attributes. For more information about the Oracle Spatial network data model, see references 1,2,3.

In Oracle Spatial 10g, NDM APIs use an in-memory approach to analyzing networks. The entire network is loaded into memory from the database. Once the network is loaded in memory, users can query and edit it. This approach works well for networks that can be completely loaded into memory; however, it cannot handle networks that are too big to be fit in memory. To address this scalability issue, Oracle Spatial 11g introduces a load-on-demand (LOD) approach to handle large networks. Instead of loading the whole network into memory, the network is first divided into manageable subnetworks (network partitions), and only partitions that are needed during analysis are loaded into memory. Loading and unloading of network partitions are automatically managed, thus removing memory as a limiting factor.

Figure 1 shows a US major highway road network. The complete US road network contains around 20 million nodes and 50 millions links (source: NAVTEQ 2006).

Networks of this size cannot be handled in-memory and need to be handled by a load-on-demand approach.




Figure 1. US road network (20 million nodes and 50 millions links, source: NAVTEQ)

LOD uses the same network data model in database as the in-memory approach: network node, link, and path tables, and the network metadata. In addition, though, LOD requires networks to be partitioned first. NDM provides partitioning procedures for spatial and logical networks, which enable users to partition their networks into network partitions and store the result in a partition table. To further speed up partition loading, this procedure can also create a BLOB representation for each network partition in a partition BLOB table.

LOD provides the following analysis functions: shortest path and hierarchical shortest path, nearest neighbors, within-cost, reachable and reaching nodes, shortest multiple-stop route (traveling salesman route), k shortest paths, network buffers, and within-cost polygon. Additional features such as network constraints, multiple cost support, partial-link paths, and user-defined data are also supported.

This paper is organized as follows: it presents the LOD network data model schema, APIs, and architecture; shows how to use the data model; discusses how to use features

such as network constraints, multiple link costs, a dynamic data set, and points along links to enhance modeling and analysis capabilities; compares major differences between the LOD and in-memory APIs; explains new features in Oracle Spatial 11g Release 2; and describes a tool for visualizing LOD networks and analysis results.

Oracle Spatial Network Data Model

The network data model consists of two parts: a network schema and network APIs. The network schema is the persistent data storage for storing network information. LOD uses the same NDM tables but with additional tables for partitions. The network APIs contain a PL/SQL package for data management in the database, a Java API for data management, and analysis on the client side or middle tier.

Network Data Model Schema

A network contains network metadata that includes network tables (node, link, path, and subpath tables). User-defined data can also be managed by NDM in Oracle Spatial 11g. In addition to the NDM network schema, LOD requires additional partition tables. Figure 2 shows the NDM schema.




Figure 2. Oracle Spatial Network Data Model (Schematic View)

Network Metadata

Network metadata provides general information about networks. Partition information needs to be inserted into the network metadata after a network is partitioned.

Network Tables

An Oracle Spatial network contains at least a node table and a link table, and a path table can be added if needed. Figure 2 shows the schema for the network data model, which includes these tables. The schema represents the information necessary for network management and analysis. User-defined data (application attributes) can be added to these tables. Node views and link views are also supported.

The network data model can also handle geometry information. That is, the network data model can represent both logical and spatial network applications. Adding geometric data to a logical network will allow the logical network to be displayed.

User-Defined Data

In Oracle Spatial 11g, users can define their own data in the user data metadata view, and NDM will manage the user data as well as the connectivity information. User-defined data is commonly used to include application attributes in the NDM representation. Complex application logic can then be implemented as network constraints or cost calculators with user-defined data.

Users can categorize different user data into different categories. Category 0 should be reserved for user data that is shared among all applications on this network. For example, in a road network, speed limit on a link belongs to category 0, because it can be used by any applications to compute the traveling time on the link. Trucking-related attributes, such as maximum truck weight, height, or length, should be assigned a category ID greater than 0, such as category 1; while hazard zone related attributes should belong to another category ID greater than 0, such as category 2.

Reading and writing of category 0 user data is managed by NDM automatically. When partition BLOBs are generated, category 0 user data is automatically included in the partition BLOBs. Users must implement their own LODUserDataIO interface for categories greater than 0 in order to load that user data into a partition.

Each specified user-defined data contains the following information:

- DATA_NAME: the name of the user data. It corresponds to the column name of the data in the node, link, or path table.
- TABLE_TYPE: the table type of the user data {NODE, LINK, PATH, SPATH}
- DATA_TYPE: the data type of the user data {VARCHAR2, NUMBER, INTEGER, SDO_Geometry}. The corresponding data types in Java are: {String, Double, Integer, JGeometry}.
- DATA_LENGTH: length for VARCHAR2 user data
- CATEGORY_ID: ID that represents the user data category

The following example shows how to insert link user data into the metadata:

```
-- Insert link user data named 'interaction' of type varchar2 (50) in
-- network 'bi_test'.
-- 'interaction' is a column of type varchar2(50) in the link table of
-- network 'bi_test'.
insert into user_sdo_network_user_data
(network,table_type,data_name,data_type,data_length,category_id)
values ('bi_test','LINK','interaction','VARCHAR2',50, 0) ;
-- Insert link user data named 'PROB' of type Number.
-- 'PROB' is a column of type NUMBER in the link table of network 'bi_test'.
insert into user_sdo_network_user_data(network,table_type,data_name,data_type,category_id)
values ('bi_test','LINK','PROB','NUMBER', 0) ;
```

Once a network or network partition is loaded, user-defined data is available in NDM Java representations. You can access user-defined data through the `getUserData` and `setUserData` methods for the Node, Link, Path, and SubPath interfaces. For example:

```
String interaction = (String)link.getUserData("interaction");
double prob = ((Double)link.getUserData("PROB")).doubleValue();
```

Partition Tables

LOD requires networks to be partitioned. Once a network is partitioned, the partition result is stored in the partition tables, which are described as follows:

Network Partition Table

- `NODE_ID` (NUMBER): network node ID
- `PARTITION_ID` (NUMBER): network partition ID
- `LINK_LEVEL` (NUMBER): network link level

Network Partition BLOB Table

- `PARTITION_ID` (NUMBER): network partition ID
- `LINK_LEVEL` (NUMBER): network link level
- `BLOB` (BLOB): partition BLOB
- `NUM_INODES` (NUMBER): number of internal partition nodes
- `NUM_ENODES` (NUMBER): number of external partition nodes
- `NUM_ILINKS` (NUMBER): number of internal partition links
- `NUM_ELINKS` (NUMBER): number of boundary partition links
- `NUM_INLINKS` (NUMBER): number of incoming partition links
- `NUM_OUTLINKS` (NUMBER): number of outgoing partition links
- `USER_DATA_INCLUDED` (VARCHAR2(1)): 'Y' if the partition contains user-defined data, 'N' otherwise

The partition table name and partition BLOB table name are stored in network metadata.

Network Partitioning

The NDM node-based partition approach divides network nodes into partitions. Each node is associated with a partition ID. In addition, NDM considers link priorities, so that link priority levels (link levels) can be considered. Each link is assigned a priority to its `link_level` column. Link priorities have positive integer values starting from 1.

For example, in road networks links are assigned with priorities based on their speed limits. One way to represent such link priorities is to designate local roads as `link_level` = 1 and state or interstate highways as `link_level` = 2, thus creating a two-level network. In such multilevel

networks, routing can be computed based on the target link level where primary path computation occurs. Note that lower link levels contain all nodes and links of their higher levels.

NDM provides a spatial partitioning procedure to help users partition networks. This procedure generates network partitions based on the maximum number of nodes per partition as specified by the user. A bisecting approach is performed recursively on a network until the desired number of nodes per partition is achieved. The number of partitions is always a power of 2. In the beginning there is only one big partition that contains all nodes. Each subsequent bisecting will divide the partitions into halves until the number of nodes in each partition is less than or equal to the given maximum number of nodes per partition. The partitions are of equal size for each link level.

For example, partitioning a network of 1,000,000 nodes with the maximum number of node per partitions equal to 5000 will generate 256 partitions (3907 nodes/partition). The following formula can be used to obtain the number of partitions, P , based upon the total number of nodes, T_v , and maximum number of nodes per partition, P_v :

$$P = 2^N \text{ where } N = \text{CEIL}(\text{LN}(T_v/P_v) / \text{LN}(2)), \text{ the number of bisections}$$

Note:

- CEIL returns the smallest integer value that is greater than or equal to the argument, and LN returns the natural logarithm of the argument.
- The actual nodes per partition = T_v/P_v .

The result is stored in a partition table. To further speed up partition loading, BLOB representations can be generated and stored in a partition BLOB table. This procedure automatically inserts partition information into its network metadata.

NDM LOD Architecture and APIs

LOD can be deployed in a simple client-server or a multi-tier environment. Figure 3 shows a typical 3-tiered architecture for LOD.




Figure 2. Oracle Spatial Network Data Model Architecture

LOD Network Analysis Capabilities

The following analysis functions are supported in the LOD API:

- Shortest Path: Find the shortest path from node A to node B.
- Accessibility Analysis: Is node A accessible from node B?
- Within-Cost Analysis: What nodes are within a given cost from (to) a given node?
- Nearest Neighbors: What are the N nearest neighbors of a given node?
- Connected Components Analysis: Label connected components with component IDs.
- Hierarchical Shortest Path: Find the shortest path based on link priority (link level).
- Multiple-Stop Shortest Path (TSP Route): Find the shortest path that passes through a given set of nodes exactly once.
- K shortest paths: Find the K shortest paths from node A to node B.

- Network Buffer: Buffer (nodes, links and partial links) within a given cost from a center.
- Within-Cost polygon: Find the region (polygon) that a given point can reach or can be reached within a given cost.
- Minimum Cost Spanning Tree: Find the minimum cost spanning tree on the network.

Java Representations of LOD Network Elements

The Java network representations (network, nodes, links, paths, and subpaths) are defined as Java interfaces and can therefore be extended. These interfaces specify the necessary behaviors for the network and its elements.

The following are some common LOD Java interfaces in the package `oracle.spatial.network.lod`:

- `LogicalNode`, `LogicalLink`, `LogicalPath`, and `LogicalSubPath`
Representations of logical node, link, path, and subpath
- `SpatialNode`, `SpatialLink`, `SpatialPath`, and `SpatialSubPath`
Representations of node, link, path, and subpath with spatial information (SDO_GEOmetry)
- `LogicalPartition` and `SpatialPartition`
Representations of network partitions
- `LODNetworkManager`
For getting the `NetworkIO`, `NetworkAnalyst`, etc.
- `NetworkIO`
For reading network information from the database
- `NetworkAnalyst`
For handling network analysis
- `LODNetworkConstraint` and `LODAnalysisInfo`
Representations of network constraint and analysis information
- `PointOnNet`
Representation of a point along a link with a percentage from the start node
- `LinkCostCalculator` and `NodeCostCalculator`
Representations of user-implemented costs for links and nodes
- `TspPath`
Representation of a multiple-stop path
- `NetworkBuffer`
Representations of buffer on a network

For information about each interface, see the NDM LOD Javadoc.

Software Requirements

The load-on-demand feature in the network data model is shipped with Oracle Spatial Release 11g. The PL/SQL package is pre-loaded in the database and required Java .jar files are provided; the Java API supports JDK (or JRE) version 1.5 or later. The web-based viewer is included as a demo. For more information, see the *Oracle Spatial Topology and Network Data Models* manual.

Using LOD in Network Data Model

This section explains how to use the LOD approach. There are four major steps: network creation, network partition, partition cache configuration, and load-on-demand analysis.

Creating a Network

Create the network in the database by creating and populating node, link, and (optionally) path tables.

Partitioning a Network

Partition the network by using a partitioning procedure, specifying the maximum number of nodes in each partition. The partition result is stored in a partition table, which is automatically generated, and partition metadata information is inserted into the network metadata. To enhance the performance of network loading, you can store partitions as BLOBs in a network partition BLOB table. (Note that LOD will still work on networks that are not partitioned -- these networks will be treated as single-partition networks; however, the use of multiple partitions is recommended for LOD analysis.)

A good partition strategy is to minimize the number of links between partitions, which reduces the number of partitions that need to be loaded and the probable number of times that the same partitions need to be reloaded. Moreover, partitions that are too small require excessive loading and unloading of partitions during analysis.

The recommended maximum number of nodes per partition, assuming 1 GB of memory, is between 5,000 and 10,000. You can tune the number and see what is best for your applications, considering the available memory, type of analysis, and network size. You should also consider configuring the partition caching size.

The following PL/SQL example code shows how to partition a spatial network:

```
exec sdo_net.spatial_partition
(network->'HILLSBOROUGH_NETWORK', -- network name
partition_table_name->'HILLSBOROUGH_NETWORK_PART$', -- partition table name
max_num_nodes->5000, -- max. number of nodes per partition
log_loc->'MDDIR', -- partition log directory
log_file->'hill_part.log', --partition log file name
open_mode->'w', -- partition log file open mode
link_level->1); -- link level
```

The following example creates partition BLOBS:

```
exec sdo_net.generate_partition_blobs(
  network->'HILLSBOROUGH_NETWORK', -- network name
  link_level ->1, -- link level
  partition_blob_table_name->'HILLSBOROUGH_NETWORK_PBLOB$', -- partition BLOB table name
  includeUserdata->FALSE, --whether to include user data in partition BLOBS
  log_loc->'MYDIR', -- partition log directory
  log_file->'hill_part.log', --partition log file name
  open_mode->'a'); -- partition log file open mode
```

The preceding two examples generate the necessary partition tables for the HILLSBOROUGH_NETWORK network. After executing these examples, you can check the hill_part.log file for the current status or any errors encountered during partitioning or BLOB generation.

Similarly you can partition a logical network with the sdo_net.logical_partition function. For logical networks whose degree distribution follows a power law, the sdo_net.logical_powerlaw_partition function can be used. For more information see the *Oracle Spatial Topology and Network Data Models* manual.

Configuring the Partition Cache

Before you perform network analysis, you can configure the network partition cache to optimize performance, by modifying an XML configuration file to override the default configuration. You can specify the following for each network:

- Cache size: the maximum number of nodes in partition cache
- Partitions source: from a partition table or a partition BLOB table
- Partition BLOB translator: which partition BLOB translator class will translate the partition BLOBS into a logical partition, or vice versa
- Resident partitions: ID of partitions that will not be flushed out of the cache
- Cache flushing policy: least recently used (LRU) or frequency based

A sample configuration is as follows:

```
<?xml version="1.0" encoding="UTF-8" ?>
<LODConfigs xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.oracle.com/spatial/network/lodLODConfigs.xsd"
  xmlns="http://www.oracle.com/spatial/network/lod">
  <!-- default configuration for networks not configured -->
  <defaultLODConfig>
    <LODConfig>
      <readPartitionFromBlob>false</readPartitionFromBlob>
      <partitionBlobTranslator>
        oracle.spatial.network.lod.PartitionBlobTranslator11gR2
      </partitionBlobTranslator>
      <userDataIO>oracle.spatial.network.lod.LODUserDataOSDO</userDataIO>
      <cachingPolicy>
```

```
<linkLevelCachingPolicy>
<linkLevel>1</linkLevel>
<maxNodes>500000</maxNodes>
<residentPartitions>-1</residentPartitions>
<flushRule>oracle.spatial.network.lod.LRUCachingHandler</flushRule>
</linkLevelCachingPolicy>
</cachingPolicy>
</LODConfig>
</defaultLODConfig>
...
<!-- network to be configured -->
<networkLODConfig>
<networkName>HILLSBOROUGH_NETWORK</networkName>
<LODConfig>
<!-- read partitions from partition table or from partition BLOB table -->
<readPartitionFromBlob>true</readPartitionFromBlob>
<partitionBlobTranslator>
  oracle.spatial.network.lod.PartitionBlobTranslator11gR2
</partitionBlobTranslator>
<userDataIO>oracle.spatial.network.lod.LODUserDataOSDO</userDataIO>
<cachingPolicy>
<linkLevelCachingPolicy>
<linkLevel>1</linkLevel>
<!-- Maximum number of nodes allowed in cache -->
<maxNodes>500000</maxNodes>
<!-- resident partitions -->
<residentPartitions>-1</residentPartitions>
<flushRule>oracle.spatial.network.lod.LRUCachingHandler</flushRule>
</linkLevelCachingPolicy>
<linkLevelCachingPolicy>
<linkLevel>2</linkLevel>
<maxNodes>500000</maxNodes>
<residentPartitions>*</residentPartitions>
<flushRule>oracle.spatial.network.lod.LRUCachingHandler</flushRule>
</linkLevelCachingPolicy>
</cachingPolicy>
</LODConfig>
</networkLODConfig>
</LODConfigs>
```

The configuration can be reloaded, if necessary.

LOD provides functions to help you estimate the best caching size. The recommendation, however, is based on logical network information; so if additional user-defined information is stored in network partitions, caching size should be reduced accordingly.

The following example returns the estimated size in bytes for a given network partition:

```
SELECT SDO_NET.GET_PARTITION_SIZE (
  NETWORK->'HILLSBOROUGH_NETWORK',
  PARTITION_ID->1,
```

```
LINK_LEVEL ->1,  
INCLUDE_USER_DATA->false,  
INCLUDE_SPATIAL_DATA->'Y') FROM DUAL;
```

Analyzing the Network

After you have created and partitioned the network, and optionally configured the partition cache, you can issue analysis queries. Analysis results are returned in Java representation. See the LOD Javadoc for details.

Java API (`oracle.spatial.network.lod`)

The following sample code uses the LOD Java API to issue a shortest-path query on a network:

```
// load LOD Configuration (XML)  
InputStream config = LODTest.class.getResourceAsStream(configXmlFile);  
LODNetworkManager.getConfigManager().loadConfig(config);  
// get database connection  
Connection conn = LODNetworkManager.getConnection(  
    dbUrl, dbUser, dbPassword);  
// get LOD network IO Adapter  
String networkName = "HILLSBOROUGH_NETWORK";  
NetworkIO reader = LODNetworkManager.getCachedNetworkIO(  
    conn, networkName, networkName, null);  
// get analysis module  
NetworkAnalyst analyst = LODNetworkManager.getNetworkAnalyst(reader);  
// compute the shortest path  
LogicalSubPath path = analyst.shortestPathDijkstra(  
    new PointOnNet(startNodeId),  
    new PointOnNet(endNodeId), null);  
...
```

Modeling and Analysis Enhancements

Several features enhance the flexibility of modeling and analysis, and greatly simplify the customization of application requirements. These features are described in this section.

Network Constraints

A network constraint is a user-implemented interface to interact with and guide the LOD analysis engine, based on application information and logic. The interface contains a simple Boolean function that must be implemented by users. Analysis information is passed to this implementation to help determine the feasible links and nodes. The constraint implementation also requires methods that indicate if the constraint requires complete path information and user-defined data.

The following demonstrates how to implement the LODNetworkConstraint interface.

```
public class ProhibitedTurnConstraint implements LODNetworkConstraint
{
    // prohibited turns information
    private HashMap<Long, long[]> pTurnMap = null;

    // construct prohibited turns information
    public ProhibitedTurnConstraint() {
        pTurnMap = new HashMap();
        // add prohibited turns as start link : prohibited end links
        long [] startLinkIDs = {145485662, 145483032};
        long [][] prohibitedEndLinkIDs = {{145477663}, {145477866}};
        for ( int i = 0; i < startLinkIDs.length ;i++)
            pTurnMap.put(startLinkIDs[i], prohibitedEndLinkIDs[i]);
    }

    // check if the given turn (start link ID, end link ID) is allowed
    private boolean validTurn(long startLinkID, long endLinkID) {
        if ( pTurnMap == null ) // no prohibited turns
            return true;
        else {
            long [] prohibitedEndLinks = pTurnMap.get(startLinkID);
            if ( prohibitedEndLinks == null )
                return true;
            else {
                for ( int i = 0; i < prohibitedEndLinks.length;i++) {
                    if ( prohibitedEndLinks[i] == endLinkID)
                        return false; // prohibited turn found
                }
                return true; // OK
            }
        }
    }
}
```

```
public boolean isSatisfied(LODAnalysisInfo info) {
    LogicalLink currentLink = info.getCurrentLink();
    if ( currentLink == null )
        return true; // start node, current link == null
    LogicalLink nextLink = info.getNextLink();
    return validTurn(currentLink.getId(), nextLink.getId());
}

public boolean requiresUserData(){ return false; }
}
```

Multiple Cost Support

In Oracle Spatial 11g, you can specify multiple costs on a node or a link, whereas in Oracle Spatial 10g you can associate only one set of costs (major cost) with links and nodes. The ability to specify multiple costs, combined with the use of user-defined data, enables you to switch from different costs in analysis. For example, travel time and travel distance are two commonly used costs for route computation in road networks. In this scenario, you would simply implement a method that returns a specific cost for a node or link and use it in an analysis function. The interface has the following method:

```
public double getLinkCost((LODAnalysisInfo analysisInfo);
```

With such an implementation, the user-specified cost will override the default cost specified in network metadata. This interface can be applied on node cost as well.

The following shows how users can implement their own link cost function for use in cost-related analysis. In this example, the default link cost is the value of the cost column of the link table, and the user-defined link cost is the travel time computed from the default link cost and the speed limit for the link level.

```
public class TravelTimeCalculator implements LinkCostCalculator
{
    double SPEED_LIMIT_1 = 30;
    double SPEED_LIMIT_2 = 60;

    public double getLinkCost(LODAnalysisInfo analysisInfo)
    {
        LogicalLink link = analysisInfo.getNextLink(); // find the link to be traversed
        int linkLevel= link.getLevel();
        double linkCost = link.getCost(); // find the distance
        switch(linkLevel) // return the travel time as distance/speed limit
        {
            case 2:
                return linkCost/SPEED_LIMIT_2;
            default:
                return linkCost/SPEED_LIMIT_1;
        }
    }
}
```

```
public int[] getUserDataCategories()
{
    return null;
}

public static void main(String[] args)
{
    ...
    ...
    //get network analyst
    NetworkAnalyst analyst=LODNetworkManager.getNetworkAnalyst(reader);

    //Shortest path analysis using default cost
    LogicalPath path =
        analyst.shortestPathDijkstra(new PointOnNet(startNodeId),
        new PointOnNet(endNodeId), null);

    //Set travel time as cost
    LinkCostCalculator lcc = new TravelTimeCalculator();
    analyst.setLinkCostCalculator(lcc);

    //Shortest path analysis using travel time as cost
    LogicalPath path =
        analyst.shortestPathDijkstra(new PointOnNet(startNodeId),
        new PointOnNet(endNodeId), null);
    ...
}
```

Precomputed Connected Components

Checking if two nodes are connected is a common query in network applications. You can use this information to eliminate unnecessary computations if two nodes are not connected. In Oracle Spatial 11g, you can precompute connected component information and store it in a connected component table. Each node at the various link levels is given a component ID.

Once the connected component table is generated, you can quickly decide if two nodes are connected by checking their component IDs. (If node A can reach node B, nodes A and B are considered to be connected.) If two nodes are not connected, there is no possible path between node A and B. This information can be used as a filter to avoid unnecessary path computations.

The connected component table is described as follows:

- NODE_ID (NUMBER): network node ID
- COMPONENT_ID (NUMBER): network component ID
- LINK_LEVEL (NUMBER): network link level

Note that the information in connected component table is precomputed and needs to be recomputed if network connectivity is changed.

The PL/SQL procedure for generating the connected components of a network is shown in the following example:

```
exec sdo_net.find_connected_components(
  network->'HILLSBOROUGH_NETWORK', -- network name           |
  link_level->1, -- target link level
  component_table_name->'NYC_COMP$', -- component result table
  log_loc->'MYDIR', -- partition log directory
  log_file->'nyc.log', -- partition log file name
  open_mode ->'a'); -- partition file mode
```

This procedure automatically inserts the connected component table name in the network metadata.

Dynamic Data Set

The dynamic data set is network data (not metadata) that users can temporarily modify in order to affect network analysis. For example, if you want to disable some links in the network temporarily (such as to indicate road segments closed for repairs or by bad weather), you can put the disabled link information in the network dynamic data set and pass the dynamic data set to the analysis function. In this case, for example, when a shortest path is computed, any temporarily disabled links will not be included in the path.

The dynamic data set contains element state, including any changes that need to be reflected in subsequent analysis requests. Users first obtain network elements from the database, but they can then modify attributes of these elements. During network analysis, the dynamic data set overrides any matching element in the data originally loaded from the network partitions. The dynamic data set can be used on a per-query basis or be reused for multiple queries.

The following example shows how the dynamic data set is used to make some links inactive during analysis:

```
...
//get network input/output object
NetworkIO reader = LODNetworkManager.getCachedNetworkIO(
  conn, networkName, networkName);

//construct dynamic data set
NetworkUpdate networkUpdate = new NetworkUpdate();
LogicalLink oldLink = reader.readLogicalLink(linkId, true);
LogicalLink newLink = (LogicalLink) oldLink.clone();
newLink.setIsActive(false);
int pid = reader.readPartitionId(newLink.getStartNodeId(), 1);
networkUpdate.updateLink(newLink, pid);
pid = reader.readPartitionId(newLink.getEndNodeId(), 1);
networkUpdate.updateLink(newLink, pid);

//get network analyst
NetworkAnalyst analyst = LODNetworkManager.getNetworkAnalyst(reader);

//Set network update in the analyst
HashMap<Integer, NetworkUpdate> networkUpdates =
  new HashMap<Integer, NetworkUpdate>();
```

```

networkUpdates.put(1, networkUpdate);
analyst.setNetworkUpdate(networkUpdates);

//find the shortest path with updated network information
LogicalSubPath path = analyst.shortestPathDijkstra(
    new PointOnNet(startNodeID), new PointOnNet(endNodeID), null);

...

```

Points Along A Link

Most analysis functions deal with network nodes as query inputs. Some applications require targets to be points on a link. For instance, a house address is usually represented as a road segment and a percentage from the start of the road segment. NDM introduces a representation called PointOnNet for this purpose. All NDM analysis functions can take PointOnNet as well as nodes as inputs. The following example creates a middle point along a link:

```
PointOnNet pt = New PointOnNet(linkID,0.5) ;...
```

The PointOnNet can also represent a node, as follows:

```
PointOnNet pt = New PointOnNet(nodeID);
```

```
...
```

Analysis results are usually represented as subpaths (described below) if the analysis involves points along links.

Partial Link Paths (Subpaths)

Oracle Spatial 11g introduces the partial link path (subpath) feature, to represent a part of a reference path in which the start and end points are specified as, in each case, the start node of a link and optionally a percentage of the distance along the link. A subpath refers to an existing path by the following parameters:

- Reference path ID: the path ID of the reference path
- Start link index: the start link index on the reference path
- Start percentage: the percentage of the subpath's start node on the start link
- End Link index: the end link index on the reference path
- End percentage: the percentage of the subpath's end node on the start link

The cost and geometry information is computed based on the percentage information.

Figure 4 shows a subpath. It starts at 50% (mid-point) of the first link of the reference path and ends at 50% of the last link of the reference path.




Figure 4. Subpath Representation

Hierarchical Shortest Path Computation

Shortest path computation could be expensive if the network is large and the path is long. In some applications, sub-optimal solutions are acceptable if the performance (computation time) is good. One common approach to trading path quality for performance is hierarchical path computation.

For hierarchical path computation, a network is divided into different levels based on link priorities. Links of higher priority will be traversed first. This approach dramatically reduces the search space, and usually the results are good compared to the optimal solution. A typical example is the route computation on a road network. A road network is classified into local roads (link level 1), and state and interstate highways (link level 2). The speed limit indicates link priority for computing travel time.

Figure 5 shows a two-level network. If the start and end nodes (denoted by S_n and E_n) are on link level 1, the hierarchical path will first compute the nearest node (HS_n) on link level 2 from the start node (S_n) and the nearest node (HE_n) of link level 2 to the end node (E_n). The hierarchical path is just a path from $SP(S_n \rightarrow HS_n) + SP(HS_n \rightarrow HE_n) + SP(HE_n \rightarrow E_n)$, where $SP(A \rightarrow B)$ is the shortest path from A to B. Note that hierarchical paths always try to traverse links of the target link level if possible.




Figure 5. Hierarchical Shortest Path from S_n to E_n

The following assumptions apply for hierarchical path computation:

- The network is classified into different hierarchical link levels.
- Link level i, L_i needs to connect to link level $L(i-1)$ and $L(i+1)$ if it exists.
- A hierarchical path will always try to traverse links of the target link level.

The following example shows how to compute a hierarchical shortest path on link level 2:

```
...
int linkLevel = 2;
//Compute shortest path using the preferred link level
LogicalPath path = analyst.shortestPathDijkstra(
    new PointOnNet(startNodeld), new PointOnNet(endNodeld), linkLevel, null);
...
...
```

Shortest Multiple-Stop Path

This analysis function computes the shortest path that passes through a set of given nodes exactly once. It is commonly known as the traveling salesman problem (TSP). NDM supports open and closed TSP paths. In an open TSP analysis, the start and/or end point can be fixed or unfixed. In a closed TSP analysis, the tour starts from the first input point and goes back to the first input point in the end. The result is a TspPath representation that includes a set of LogicalPath for paths between two consecutive nodes and their visit order.

The following example shows how to compute a closed shortest multiple-stop path:

```
...
PointOnNet[][] tspNodes = nodeIdsToPoints(tspNodeIds); // convert nodes to point on nets
TSP.TourFlag tourFlag = TSP.TourFlag.CLOSED; // a closed TSP tour
TspPath tspPath =
    analyst.tsp(tspNodes, tourFlag, linkLevel, constraint);
...
```

K Shortest Paths

This analysis function computes k shortest paths of a given start and end node. The following example shows how to find 3 shortest paths from a start node to an end node:

```
...
PointOnNet[] startPoint = {new PointOnNet(startNodeId)};
PointOnNet[] endPoint = {new PointOnNet(endNodeId)};
Int k = 3;
LogicalSubPath[] paths =
    analyst.kShortestPaths(startPoint, endPoint, k, constraint, null);
...
```

Network Buffer

A spatial buffer of a geometry is a geometric representation of which the boundary has the same offset (distance) to the target geometry. Similar to a spatial buffer, a network buffer is defined on a network with centers (target points) and a radius (network cost). A network buffer consists of the network elements that are within a given radius from or to the centers. A network element can be a node, a link, or part of a link. Any point covered by the network buffer has a minimum cost to the centers less than or equal to the given radius.

A network buffer consists of coverage and cost information. The centers (target points) of a network buffer are the input points on the network from which the buffer is computed. The coverage of a network buffer is represented by the set of network elements within the buffer. The cost or cost range of a network element is the minimum cost or cost range from (or to) the element to (or from) the centers. The following example shows how to construct a network buffer:

```
PointOnNet [] centerPoints = {new PointOnNet(nodeID)}; // buffer center node
double radius = 1000.0; // buffer radius
LODNConstraint constraint = null;
```

```
NetworkBuffer netBuffer = analyst.NetworkBuffer(centerPoints, radius, constraint);
...
```

NetworkBuffer is a Java in-memory representation. It provides methods to return the coverage information and cost associated with coverage. Users can persist the desired results if needed. Network buffers can be computed on spatial and logical networks as only connectivity and cost information is used. Users can make network buffers persistent for further queries.

Within-Cost Polygon

This analysis function computes the polygon that represents the region of a node can reach or be reached within a given cost. This spatial representation is usually used as a zone of influence in spatial queries. For instance, it can represent the region that a retail store can reach within 10 kilometers or the region that a hospital that can reach within 10 minutes. The following example shows how to compute the drive time polygon (with travel time as cost):

```
...
PointOnNet [] startPoints = {new PointOnNet(nodeID)};// target node
double cost = 10.0;// cost = 10 minutes
LODNetworkConstraint constraint = null;
LODGoalNodeFilter goalNodeFilter = null;
JGeometry wcPolygon = analyst.withinCostPolygon(startPoints, cost, constraint, goalNodeFilter); // drive time polygon in
JGeometry
```

Minimum Cost Spanning Tree

This analysis computes a minimum cost spanning tree on a network. The resulting spanning tree is represented by the link IDs included in this tree. The following example shows how to perform such analysis.

```
...
long[] treeLinks = analyst.mcst(linkLevel, null, null);
...
```

Comparisons Between NDM LOD and In-Memory APIs

This section explains the differences between the LOD API (introduced in Oracle Spatial 11g) and the in-memory API (using a Java API and its associated SDO_NET_MEM PL/SQL package) . It compares their approaches, data model, analysis functions, and future support:

Approaches

The LOD and in-memory APIs take the following basic approaches:

- The LOD API requires that the network be partitioned first, and then partitions are loaded into memory whenever needed during LOD analysis. The LOD API can be found under the Java package `oracle.spatial.network.lod`
- The in-memory API requires that the entire network be loaded into memory before analysis can be conducted. The in-memory API can be found under the Java package

`oracle.spatial.network`. A PL/SQL interface (SDO_NET_MEM) is built on top of the in-memory API.

The LOD and in-memory APIs cannot be mixed in an application. Choose the appropriate API according to your application requirements and network size.

Data Model

The LOD and in-memory APIs use the same network data model, but the LOD API requires additional partition information. (Although the LOD API can handle networks without partition information, such networks will be treated as networks with a single partition.)

- The LOD API can only perform read-only operations on a network. Network updates need to be done at the database level, and partitions (and partition BLOBs, if used) need to be updated accordingly. Access to network elements needs to be done through the LOD API.
- The in-memory API can perform read, write, and edit operations on a network. All network elements are available after a network is loaded.

Analysis Functions

Both the in-memory and LOD APIs support shortest path, nearest neighbors, within-cost, reachable and reaching analysis, minimum-cost spanning, shortest multiple-stop path (TSP), and k-shortest paths operations, as well as network constraints and user-defined data. However, LOD also supports hierarchical shortest-path computation, network buffers, and within-cost polygon computation.

The LODNetworkManager class is the main entry for LOD analysis, and the NetworkManager class is the main entry for in-memory analysis. Supported analysis functions can be found in the Javadoc for each class.

Future Support

The in-memory APIs are deprecated in future releases due to its scalability limitation on handling large networks. All modeling features and analysis functions in the in-memory API are supported in the LOD API. In addition, new features and analysis functions have been added to the LOD API.

We believe that the LOD approach provides a flexible, scalable, and high-performance network modeling and analysis platform for many network applications. Users are strongly encouraged to adopt the LOD approach for their applications. The Oracle Spatial 11g Release 2 routing engine is now using the NDM LOD API as its network analysis engine. The NDM LOD API provides a data translator that directly supports routing network data (in NAVTEQ ODF format). For more information on ODF, see reference 5.

What Is New in Oracle Spatial 11g Release 2

The following is a list of new Oracle Spatial 11g Release 2 features for the NDM LOD API:

- Modeling Features
 - PointOnNet
 - LinkCostCalculator and NodeCostCalculator

Note that the new signatures take LODAnalysisInfo as inputs, while the 11gR1 signatures took LogicalLink and LogicalNode as inputs. The new definition is more general and can model much complex costs of multiple links or nodes.

- NetworkBuffer
- Analysis Functions
 - K shortest paths function
 - Shortest multiple stop path (TSP) function
 - Network buffer function
 - Within-cost polygon function
 - Logical partition functions (sdo_net.logical_partition, sdo_net.logical_powerlaw_partition)

Web-Based LOD Analysis Demo

A web-based demo is provided to show users how to visualize LOD analysis results. This demo is implemented in a 3-tiered architecture with the JavaScript GUI (web browser) built with the Oracle MapViewer JavaScript API to handle user interactions, the NDM analysis engine deployed in the middle tier, and the network data stored in the database backend.

To install the demo, first download sample network data from NAVTEQ (see 5). Once the sample network is populated and registered in the database, deploy the ndmdemo.ear file in your web server.

To run the demo, first enter the configuration information in the getConfiguration.jsp page, where you can specify the database, mapviewer and geocoder connection information. You need to have a MapViewer service running on the same machine as your NDM servlet. For information on how to install and use MapViewer, see the *Oracle Fusion Middleware User's Guide to Oracle MapViewer* [4].

After configuring the demo, you can perform various analysis operations on the selected network and use Oracle MapViewer functions for visualizing analysis results. A text area prints out analysis result and statistics. Figure 6 shows the home page of the demo. Figure 7 shows a shortest path analysis result.

Oracle Spatial Network Data Model Demo - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://stadm68.us.oracle.com:7001/ndmdemo/index.html

Oracle Spatial Network Data Model Demo

This web application shows how to use NDM, router, geocoder and mapviewer to conduct network analysis and draw results on a map. The purpose of this demo is to present the usage of NDM together with other spatial components in a simple and readable way. The source code in this demo is by no means the best practice. Users of NDM should organize their programs according to the requirements of their applications.

Before running this demo, you need to:

- set up the ndmdemo and router_sd data. To set up the data, download the [data.zip](#) file and follow the README in the package.
- deploy and configure geocoder. The geocoder should connect to the router_sd user in your database.
- deploy and configure mapviewer. You need to create a mapviewer data source named "ndmdemo" connecting to the ndmdemo user in your database.

Before accessing the network analysis pages, you must first set the geocoder and database connections using [getConfiguration.jsp](#).

Demo	Source Code	Description
getConfiguration.jsp	getConfiguration.jsp setConfiguration.jsp Configuration.java	Configuration Information
shortestPath.jsp	shortestPath.jsp shortestPathAnalysis.jsp	Find the shortest path from the start location to the end location. This demo also shows how to apply custom constraints and link cost calculators.
nearestNeighbors.jsp	nearestNeighbors.jsp nearestNeighborsAnalysis.jsp	Find the nearest nodes from the start location
withinCost.jsp	withinCost.jsp withinCostAnalysis.jsp	Find nodes within the given cost
withinCostPolygon.jsp	withinCostPolygon.jsp withinCostPolygonAnalysis.jsp	Find the polygon covering the area within the given cost
networkBuffer.jsp	networkBuffer.jsp networkBufferAnalysis.jsp	Find the network buffer within the give cost
kShortestPaths.jsp	kShortestPaths.jsp kShortestPathsAnalysis.jsp	Find K shortest paths
tsp.jsp	tsp.jsp tspAnalysis.jsp	Traveling Salesman Problem Analysis
showMap.js	showMap.js	JavaScript functions to render the analysis results on a map
AnalysisEngine.java	AnalysisEngine.java	Java class that processes analysis requests
Constraint Implementations	NoHighwayConstraint.java ProhibitedZoneConstraint.java	Java classes that implement LODNetworkConstraint
Link Cost Calculator Implementations	LinkHopCostCalculator.java LinkTimeCostCalculator.java	Java classes that implement LinkCostCalculator
Goal Node Filter Implementations	EvenNodeFilter.java	Java classes that implement LODGoalNode

Figure 6. NDM Web-Based Demo – Home Page

NDM Demo: Shortest Path Analysis - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://stadm68.us.oracle.com:7001/ndmdemo/shortestPath.jsp

Oracle Spatial Network Data Model Demo

Shortest Path Analysis

Left click for start point, right click for end point, or manually enter node ID, link ID@percentage, or address.

Start

End

Network Constraints (Multiple Select)

custom.NoHighwayConstraint
oracle.spatial.router.ndm.TruckHeightConstraint
oracle.spatial.router.ndm.TruckLegalConstraint
oracle.spatial.router.ndm.TruckLengthConstraint

Link Cost Calculators

Default Link Cost Calculator

Keep Previous Results
Reverse Direction

Analysis Result:

[0]:205251671->205031635)
[cost:14880.68995, 114 links]

Figure 7. NDM Web-Based Demo – Shortest Path Analysis

Conclusion

The load-on-demand feature in the Oracle Spatial network data model, available in Oracle Spatial Release 11g, is designed to handle analysis operations on large networks. It provides a scalable and a flexible solution to many network applications. We are currently working with our customers and partners to extend the modeling and analysis capabilities of the network data model.

For more technical information about the network data model, see the Oracle Spatial web site:
<http://www.oracle.com/technology/products/spatial/index.html>.

References

1. *Oracle 10g, Oracle Spatial Network Data Model, A Technical White Paper*, Oracle Corporation, May 2005
2. *Building GIS Applications Using the Oracle Spatial Network Data Model: A Technical White Paper*, Oracle Corporation, May 2005
3. *Oracle Spatial Topology and Network Data Models*, Oracle Corporation.
4. *Oracle Fusion Middleware User's Guide for Oracle MapViewer* (previously titled *OracleAS MapViewer User's Guide*), Oracle Corporation
5. NAVTEQ, Oracle Delivery Format (ODF), <http://developer.navteq.com>



A Load-On-Demand Approach to Handling
Large Networks in the Oracle Spatial Network

Data Model

February 2011

Authors: Jack Chenghua Wang

Contributing Author: Huijing Gong, Betsy
George

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

Hardware and Software, Engineered to Work Together