

# Oracle8™ Spatial Cartridge

User's Guide and Reference

Release 8.0.4

November 1997

Part No. A53264-02

---

Oracle8 Spatial Cartridge User's Guide and Reference

Part No. A53264-02

Release 8.0.4

Copyright © 1997, Oracle Corporation. All rights reserved.

Primary Author: Jeff Hebert

Contributing Author: Anna Logan

Contributors: Dan Geringer, Peter Vretanos, Jayant Sharma

**The programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be licensee's responsibility to take all appropriate fail-safe, back up, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and Oracle disclaims liability for any damages caused by such use of the Programs.**

This Program contains proprietary information of Oracle Corporation; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright patent and other intellectual property law. Reverse engineering of the software is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error free.

If this Program is delivered to a U.S. Government Agency of the Department of Defense, then it is delivered with Restricted Rights and the following legend is applicable:

**Restricted Rights Legend** Programs delivered subject to the DOD FAR Supplement are 'commercial computer software' and use, duplication and disclosure of the Programs shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, Programs delivered subject to the Federal Acquisition Regulations are 'restricted computer software' and use, duplication and disclosure of the Programs shall be subject to the restrictions in FAR 52.227-14, Rights in Data -- General, including Alternate III (June 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

Oracle, SQL\*Forms, SQL\*Loader, SQL\*Menu, SQL\*Net, SQL\*Plus, and SQL\*Repor are registered trademarks of Oracle Corporation. Oracle7 and Oracle8 are trademarks of Oracle Corporation.

HHCODE is a trademark of the Government of Canada.

All other products or company names are used for identification purposes only, and may be trademarks of their respective owners.

---

# Contents

<b>Send Us Your Comments .....</b>	<b>vii</b>
------------------------------------	------------

<b>Preface.....</b>	<b>ix</b>
---------------------	-----------

## **1 Spatial Cartridge Concepts**

1.1	Introduction to Spatial Data.....	1-1
1.2	Geometric Types.....	1-2
1.3	Data Model.....	1-3
1.3.1	Element .....	1-3
1.3.2	Geometry .....	1-3
1.3.3	Layer.....	1-4
1.4	Database Structures.....	1-4
1.5	Indexing Methods .....	1-9
1.5.1	Tessellation of a Layer .....	1-9
1.5.2	Fixed-Size Tile Spatial Indexing.....	1-11
1.5.3	Variable-Sized Tile Spatial Indexing .....	1-14
1.6	Partitioned Point Data .....	1-17

## **2 Loading Spatial Data**

2.1	Load Model .....	2-1
2.2	Load Process.....	2-2
2.2.1	Bulk Loading.....	2-2
2.2.2	Transactional Insert Using SQL .....	2-4
2.2.3	Transactional Insert Using Spatial Geometry Functions.....	2-6
2.3	Index Creation .....	2-8

2.3.1	Choosing a Tessellation Algorithm .....	2-8
2.3.2	Spatial Indexing with Fixed-Size Tiles .....	2-9
2.3.3	Spatial Indexing with Variable-Sized Tiles .....	2-12

### 3 Querying Spatial Data

3.1	Query Model .....	3-1
3.2	Spatial Data Model .....	3-2
3.3	Spatial Query .....	3-5
3.3.1	Dynamic Query Window .....	3-6
3.3.2	Primary Filter .....	3-7
3.3.3	Secondary Filter .....	3-8
3.4	Spatial Join .....	3-9

### 4 Partitioning Point Data

4.1	Overview .....	4-1
4.2	Partitioning Process .....	4-1
4.3	Function Details .....	4-2

### 5 Administrative Procedures

SDO_ADMIN.POPULATE_INDEX .....	5-2
SDO_ADMIN.POPULATE_INDEX_FIXED .....	5-4
SDO_ADMIN.SDO_CODE_SIZE .....	5-7
SDO_ADMIN.UPDATE_INDEX .....	5-8
SDO_ADMIN.UPDATE_INDEX_FIXED .....	5-10
SDO_ADMIN.VERIFY_LAYER .....	5-12
Partitioned Point Data Procedures .....	5-13
SDO_ADMIN.ALTER_HIGH_WATER_MARK .....	5-14
SDO_ADMIN.DROP_PARTITION_INFO .....	5-15
SDO_ADMIN.PARTITION .....	5-16
SDO_ADMIN.PROPAGATE_GRANTS .....	5-18
SDO_ADMIN.REGISTER_PARTITION_INFO .....	5-19
SDO_ADMIN.REPARTITION .....	5-20

SDO_ADMIN.VERIFY_PARTITIONS .....	5-22
-----------------------------------	------

## 6 Tuning Functions

SDO_TUNE.ESTIMATE_TILING_LEVEL .....	6-2
SDO_TUNE.EXTENT_OF .....	6-5

## 7 Geometry Functions

SDO_GEOM.ADD_NODES.....	7-2
SDO_GEOM.INIT_ELEMENT .....	7-4
SDO_GEOM.INTERACT .....	7-5
SDO_GEOM.RELATE.....	7-7
SDO_GEOM.VALIDATE_GEOMETRY.....	7-10

## 8 Window Functions

SDO_WINDOW.BUILD_WINDOW .....	8-2
SDO_WINDOW.BUILD_WINDOW_FIXED.....	8-4
SDO_WINDOW.CLEAN_WINDOW.....	8-6
SDO_WINDOW.CREATE_WINDOW_LAYER .....	8-7

## 9 Partitioned Point Data Functions

SDO_BVALUETODIM .....	9-3
SDO_COMPARE.....	9-4
SDO_DATETODIM.....	9-6
SDO_DECODE.....	9-8
SDO_ENCODE .....	9-9
SDO_TO_BVALUE.....	9-10
SDO_TO_DATE.....	9-11

## A Sample SQL Scripts and Tuning Tips

A.1	Sample SQL Scripts .....	A-1
A.1.1	Scripts for Spatial Indexing.....	A-1

A.1.1.1	cr_spatial_index.sql Script .....	A-1
A.1.1.2	crlayer.sql Script .....	A-2
A.1.2	Scripts for Partitioned Point Data .....	A-2
A.1.2.1	altpart.sql Script.....	A-3
A.1.2.2	drppart.sql Script.....	A-3
A.1.2.3	sdogrant.sql Script.....	A-3
A.2	Tuning Tips .....	A-4
A.2.1	Data Modeling .....	A-4
A.2.2	Understanding the Tiling Level .....	A-4
A.2.3	Database Sizing.....	A-5
A.2.4	Tuning Point Data .....	A-6
A.2.4.1	Efficient Queries for Point Data .....	A-6
A.2.4.2	Efficient Schema for Point Layers.....	A-7
A.2.5	Tuning Spatial Join Queries.....	A-8
A.2.5.1	Using the NO_MERGE, INDEX, and USE_NL Hints.....	A-8
A.2.5.2	Spatial Join Queries with Point Layers .....	A-9
A.2.6	Using Customized Geometry Types .....	A-10
A.2.7	Performing Secondary Filter Queries and the Redo Log.....	A-11
A.2.8	Visualizing the Spatial Index (Drawing Tiles) .....	A-11

## **B Data Dictionary**

## **C Messages and Codes**

## **Glossary**

---

---

# Send Us Your Comments

## Oracle8 Spatial Cartridge User's Guide and Reference, Release 8.0.4

Part No. A53264-02

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, please indicate the chapter, section, and page number (if available).

You can send comments to us in the following ways

- e-mail: [nedc\\_doc@us.oracle.com](mailto:nedc_doc@us.oracle.com)
- FAX: 603.897.3269. Attn: Spatial Cartridge
- postal service:  
Oracle Corporation  
Oracle Spatial Cartridge Documentation  
One Oracle Drive  
Nashua, NH 03062  
USA

If you would like a reply, please include your name, address, and telephone number.

---

---





---

---

# Preface

The *Oracle Spatial Cartridge User's Guide and Reference* provides user and reference information for Spatial Cartridge, and extensions to Oracle8 Enterprise Edition.

Spatial Cartridge requires Oracle8 Enterprise Edition. Oracle8 and Oracle8 Enterprise Edition have the same basic features. However, several advanced features, such as data cartridges, are available only with the Enterprise Edition, and some of these features are optional. For example, to use Oracle8 table partitioning, you must have the Enterprise Edition and the Partitioning Option.

For information about the differences between Oracle8 and the Oracle8 Enterprise Edition and the features and options that are available to you, see *Getting to Know Oracle8*.

---

---

**Note :** Spatial Cartridge release 8.0.3 introduced two distinct algorithms for building a spatial index: fixed-size tiling and variable-sized tiling. Based on testing and customer feedback, for release 8.0.4, Oracle recommends using only fixed-size tiling on production systems. Variable-sized tiling, while it has theoretical advantages in some situations, is included primarily for experimentation purposes. In the future, Oracle may issue a technical bulletin further refining the usage of variable-sized tiling.

---

---

## Intended Audience

This manual is intended for anyone who needs to store spatial data in a relational database.

## Structure

This manual contains eight chapters, several appendixes, and a glossary:

Chapter 1	Introduces spatial data concepts.
Chapter 2	Explains spatial data loading.
Chapter 3	Explains methods for querying a spatial database.
Chapter 4	Explains how to use partitioned point data.
Chapter 5	Provides the syntax and semantics for the administrative functions and procedures.
Chapter 6	Provides the syntax and semantics for the tuning functions and procedures.
Chapter 7	Provides the syntax and semantics for the geometric functions and procedures.
Chapter 8	Provides the syntax and semantics for the window functions and procedures.
Chapter 9	Provides the syntax and semantics for functions and procedures relevant only to using partitioned point data.
Appendix A	Describes sample SQL scripts and tuning tips.
Appendix B	Describes the Spatial Cartridge data dictionary.
Appendix C	Provides a list of error messages and conditions.
Glossary	Provides definitions of terms used in this manual.

## Related Documents

For more information, see the following manuals:

- *Getting to Know Oracle8*
- *Oracle8 Administrator's Guide*
- *Oracle8 Error Messages*
- *Oracle8 Utilities*

For additional information about Spatial Cartridge, including a demo, several white papers, and other assorted collateral, visit the official Spatial Cartridge web site: <http://www.oracle.com/st/cartridges/spatial/>

## Conventions

In examples, an implied carriage return occurs at the end of each line, unless otherwise noted. You must press the Return key at the end of a line of input.

The following conventions are used in this manual:

---

Convention	Meaning
. . .	Vertical ellipsis points in an example mean that information not directly related to the example has been omitted.
...	Horizontal ellipsis points in statements or commands mean that parts of the statement or command not directly related to the example have been omitted
<b>boldface text</b>	Boldface type in text indicates a term defined in the text, the glossary, or in both locations.
< >	Angle brackets enclose user-supplied names.
[ ]	Brackets enclose optional clauses from which you can choose one or none.
%	The percent sign represents the system prompt on a UNIX system.

---



---

# Spatial Cartridge Concepts

Oracle Spatial Cartridge is an integrated set of functions and procedures that enables spatial data to be stored, accessed, and analyzed quickly and efficiently in an Oracle8 database.

Spatial data represents the essential location characteristics of real or conceptual objects as those objects relate to the real or conceptual space in which they exist.

## 1.1 Introduction to Spatial Data

Spatial Cartridge is designed to make the storage, retrieval, and manipulation of spatial data easier and more natural to users such as a Geographic Information System (GIS). Once this data is stored in an Oracle8 relational database, it can be easily and meaningfully manipulated and retrieved as it relates to all the other data stored in the database.

A common example of spatial data can be seen in a road map. A road map is a two-dimensional object that contains points, lines, and polygons that can represent cities, roads, and political boundaries such as states or provinces. A road map is a visualization of geographic information. The location of cities, roads, and political boundaries that exist on the surface of the Earth are projected onto a two-dimensional display or piece of paper, preserving the relative positions and relative distances of the rendered objects.

The data that indicates the Earth location (latitude and longitude, or height and depth) of these rendered objects is the spatial data. When the map is rendered, this spatial data is used to project the locations of the objects on a two-dimensional piece of paper. A GIS is often used to store, retrieve, and render this Earth-relative spatial data.

Other types of spatial data that can be stored using Spatial Cartridge besides GIS data include data from computer-aided design (CAD) and computer-aided manu-

facturing (CAM) systems. Instead of operating on objects on a geographic scale, CAD/CAM systems work on a smaller scale such as for an automobile engine or much smaller scale as for printed circuit boards.

The differences among these three systems are only in the scale of the data, not its complexity. They might all actually involve the same number of data points. On a geographic scale, the location of a bridge can vary by a few tenths of an inch without causing any noticeable problems to the road builders. Whereas, if the diameter of an engine's pistons are off by a few tenths of an inch, the engine will not run. A printed circuit board is likely to have many thousands of objects etched on its surface that are no bigger than the smallest detail shown on a roadbuilder's blueprints.


## 1.2 Geometric Types

Spatial Cartridge supports three geometric primitive types and geometries composed of collections of these types. The three primitive types are as follows:

- 2-D Point and Point Cluster
- 2-D Line Strings
- 2-D N-Point Polygons

**2-D points** are elements composed of two ordinates, X and Y, often corresponding to longitude and latitude. **Line strings** are composed of one or more pairs of points that define line segments. **Polygons** are composed of connected line strings that form a closed ring and the interior of the polygon is implied. Figure 1-1 illustrates the supported geometric primitive types.

*Figure 1-1 Geometric Primitive Types*



Self-crossing polygons are not supported although self-crossing line strings are. If a line string crosses itself it does not become a polygon. A self-crossing line string does not have any implied interior.

## 1.3 Data Model

The Spatial Cartridge data model is a hierarchical structure consisting of elements, geometries, and layers, which correspond to representations of spatial data. Layers are composed of geometries (or geometric objects), which in turn are made up of elements.

For example, a point might represent a building location, a line string might be a road or flight path, and a polygon could be a state, city, zoning district, or city block.

### 1.3.1 Element

An **element** is the basic building block of a geometric feature for Spatial Cartridge. The supported spatial element types are points, line strings, and polygons. For example, elements might model star constellations (point clusters), roads (line strings), and county boundaries (polygons). Each coordinate in an element is stored as an X,Y pair.

**Point data**<sup>1</sup> consists of one coordinate. **Line data** consists of two coordinates representing a line segment of the element. **Polygon data** consists of coordinate pair values, one vertex pair for each line segment of the polygon. Coordinates are defined in either a clockwise or counter-clockwise order around the polygon.

If an element spans more than one row, an incremental sequence number (starting at zero) orders the rows.

### 1.3.2 Geometry

A **geometry** is the representation of a user's spatial feature, modeled as an ordered set of primitive elements. Each geometric object is required to be uniquely identified by a numeric geometry identifier (GID), associating the object with its corresponding attribute set.

A complex geometric feature such as a polygon with holes would be stored as a sequence of polygon elements. In a multi-element polygonal geometry, all subelements are wholly contained within the outermost element, thus building a more complex geometry from simpler pieces.

For example, a geometry might describe the buildable land in a town. This could be represented as a polygon with holes where water or zoning prevents construction.

---

<sup>1</sup> Point data can also be stored in a partitioned table. See Chapter 4, "Partitioning Point Data" for details.

### 1.3.3 Layer

A **layer** is a heterogeneous collection of geometries having the same attribute set. For example, one layer in a GIS might include topographical features, while another describes population density, and a third describes the network of roads and bridges in the area (lines and points). Each layer's geometric objects and their associated spatial index are stored in the database in standard tables.

## 1.4 Database Structures

Spatial Cartridge uses four database tables to store and index spatial data. These four tables are collectively referred to as a *layer*. A template SQL script is provided to facilitate the creation of these tables. See Section A.1.1.2, "crlayer.sql Script" for details.

Table 1-1 through Table 1-4 describe the schema of a Spatial Cartridge layer.

**Table 1-1** *<layername>\_SDOLAYER*

SDO_ORDCNT	SDO_LEVEL	SDO_NUMTILES	SDO_COORDSYS
<number>	<number>	<number>	<varchar>

**Table 1-2** *<layername>\_SDODIM table or view*

SDO_DIMNUM	SDO_LB	SDO_UB	SDO_TOLERANCE	SDO_DIMNAME
<number>	<number>	<number>	<number>	<varchar>

**Table 1-3** *<layername>\_SDOGEOM table or view*

SDO_GID	SDO_ESEQ	SDO_ETYPE	SDO_SEQ	SDO_X1	SDO_Y1	...	SDO_Xn	SDO_Yn
<number>	<number>	<number>	<number>	<number>	<number>	...	<number>	<number>

**Table 1-4** *<layername>\_SDOINDEX table*

SDO_GID	SDO_CODE	SDO_MAXCODE **	SDO_GROUPCODE **	SDO_META
<number>	<raw>	<raw>	<raw>	<raw>

The SDO\_MAXCODE and SDO\_GROUPCODE columns are not required for the recommended indexing algorithm using fixed-size tiles.



The columns of each table are defined as follows:

#### <layername>\_SDOLAYER:

- **SDO\_ORDCNT** - The SDO\_ORDCNT column is the total number of ordinates per row in the <layername>\_SDOGEOM table. That is, the total number of data value columns, and not the number of points or coordinates. SDO\_ORDCNT should not be multiplied by the total number of dimensions per coordinate as it is already a total.
- **SDO\_LEVEL** - The SDO\_LEVEL column stores the number of times the layer should be tessellated during the indexing stage. Use the SDO\_TUNE.ESTIMATE\_TILING\_LEVEL() procedure to determine an appropriate tiling level for your data.
- **SDO\_NUMTILES** - The SDO\_NUMTILES column is the number of variable-sized tiles used to tessellate each object in the <layername>\_SDOGEOM table. This column must be set to NULL when using fixed-size tiles. ***This column is new for release 8.0.4.***
- **SDO\_COORDSYS** - The SDO\_COORDSYS column is optional; where you can indicate the name of the coordinate system, using a standard such as *POSC* or *OGIS*. ***This column is new for release 8.0.4.***

#### <layername>\_SDODIM:

- **SDO\_DIMNUM** - The SDO\_DIMNUM column is the dimension to which this row refers, starting with 1 and increasing.
- **SDO\_LB** - The SDO\_LB column is the lower bound of the ordinate in this dimension. For example, if the dimension is latitude, the lower bound would be -90.
- **SDO\_UB** - The SDO\_UB column is the upper bound of the ordinate in this dimension. For example, if the dimension is longitude, the upper bound would be 180.
- **SDO\_TOLERANCE** - The SDO\_TOLERANCE column is the distance two points can be apart and still be considered the same due to round-off errors. Tolerance must be greater than zero. If you want zero tolerance, enter a number such as 0.00005, where the number of zeroes to the right of the decimal point matches the precision of your data. The extra 5 will round up to your last decimal digit.
- **SDO\_DIMNAME** - The SDO\_DIMNAME column is used for the usual name applied to this dimension, such as *longitude*, *latitude*, *X* or *Y*.

**<layername>\_SDOGEOM:**

- **SDO\_GID** - The SDO\_GID column is a unique numeric identifier for each geometry in a layer.
- **SDO\_ESEQ** - The SDO\_ESEQ column enumerates each element in a geometry, that is, the Element **SEQ**uence number.
- **SDO\_ETYPE** - The SDO\_ETYPE column is the geometric primitive type of the element. For this release of Spatial Cartridge, the valid values are SDO\_GEOM.POINT\_TYPE, SDO\_GEOM.LINestring\_TYPE, or SDO\_GEOM.POLYGON\_TYPE (ETYPE values 1, 2, and 3, respectively). Setting the ETYPE to zero indicates that this element should be ignored. See Section A.2.6 for information on ETYPE 0.
- **SDO\_SEQ** - The SDO\_SEQ column records the order (the **SEQ**uence number) of each row of data making up the element.
- **SDO\_X1** - X value of the first coordinate.
- **SDO\_Y1** - Y value of the first coordinate.
- **SDO\_Xn** - X value of the Nth coordinate.
- **SDO\_Yn** - Y value of the Nth coordinate.

**<layername>\_SDOINDEX:**

- **SDO\_GID** - The SDO\_GID column is a unique numeric identifier for each geometry in a layer. This can be thought of as a foreign key back to the <layername>\_SDOGEOM table.
- **SDO\_CODE** - The SDO\_CODE column is the bit-interleaved ID of a tile that covers SDO\_GID. The number of bytes needed for the SDO\_CODE and SDO\_MAXCODE columns depends on the level used for tiling. Use the SDO\_ADMIN.SDO\_CODE\_SIZE( ) function to determine the size required for a given layer. The maximum number of bytes possible is 255.
- **SDO\_MAXCODE** - The SDO\_MAXCODE column describes a variable-sized logical tile, which is the smallest tile (with the longest tile ID) in the current quadrant. The SDO\_MAXCODE column is SDO\_CODE padded out one place farther than the longest allowable code name for this index. This column is not used for fixed-size tiles.
- **SDO\_GROUPCODE** - The SDO\_GROUPCODE column is a prefix of SDO\_CODE. It represents a variable-sized tile at level <layer-name>\_SDOLAYER.SDO\_LEVEL that contains or is equal to the tile repre-

sented by SDO\_CODE. This column is not used for fixed-size tiles. **This column is new for release 8.0.4.**


- **SDO\_META** - The SDO\_META column is not required for spatial queries. It provides information necessary to find the bounds of a tile. See Section A.2.8 for one possible use of this column. **This column is new for release 8.0.4.**

Spatial Cartridge provides stored procedures that assume the existence of the layer schema as described in this section. While layer tables may contain additional columns, they are required to contain at least the columns described in this section with the same column names and data types.

Figure 1-2 illustrates how a geometry is stored in the database using Spatial Cartridge. The geometry to be stored is a complex polygon with a hole in it.

**Figure 1-2 Complex Polygon**

Geometry 1013:



<layername>\_SDOLAYER

SDO_ORDCNT (number)
4

## &lt;layername&gt;\_SDODIM

SDO_DIMNUM (number)	SDO_LB (number)	SDO_UB (number)	SDO_TOLERANCE (number)	SDO_DIMNAME (varchar)
1	0	100	.05	X axis
2	0	100	.05	Y axis

## &lt;layername&gt;\_SDOGEOM

SDO_GID (number)	SDO_ESEQ (number)	SDO_ETYPE (number)	SDO_SEQ (number)	SDO_X1 (number)	SDO_Y1 (number)	SDO_X2 (number)	SDO_Y2 (number)
1013	0	3	0	P1(X)	P1(Y)	P2(X)	P2(Y)
1013	0	3	1	P2(X)	P2(Y)	P3(X)	P3(Y)
1013	0	3	2	P3(X)	P3(Y)	P4(X)	P4(Y)
1013	0	3	3	P4(X)	P4(Y)	P5(X)	P5(Y)
1013	0	3	4	P5(X)	P5(Y)	P6(X)	P6(Y)
1013	0	3	5	P6(X)	P6(Y)	P7(X)	P7(Y)
1013	0	3	6	P7(X)	P7(Y)	P8(X)	P8(Y)
1013	0	3	7	P8(X)	P8(Y)	P1(X)	P1(Y)
1013	1	3	0	G1(X)	G1(Y)	G2(X)	G2(Y)
1013	1	3	1	G2(X)	G2(Y)	G3(X)	G3(Y)
1013	1	3	2	G3(X)	G3(Y)	G4(X)	G4(Y)
1013	1	3	3	G4(X)	G4(Y)	G1(X)	G1(Y)

In this example, the <layername>\_SDOGEOM table is shown as an eight column table with four ordinates per row. In actual usage, Spatial Cartridge supports N-wide<sup>1</sup> tables. The coordinates for the outer polygon in this example could have been loaded into a single row containing values for coordinates P1 to P8, and then repeating P1 to close the polygon. The coordinates would be stored in the SDO\_X1 and SDO\_Y1 through SDO\_X9 and SDO\_Y9 columns.

The data in the <layername>\_SDOINDEX table is described in Section 1.5, “Indexing Methods”.

<sup>1</sup> A <layername>\_SDOGEOM table can have up to 255 columns. The maximum number of data columns is 255, minus 4 for the other required spatial columns, and minus any other user-defined columns. For polygon and line strings, storing 16 to 20 ordinates per row is suggested for performance reasons, but not required.

## 1.5 Indexing Methods

A spatial index is considered a logical index as opposed to a physical index. The entries in the spatial index are dependent on the location of the geometric objects in the layer space, and are not dependent on the stored location of the data on the disk. This means that a table containing spatial data could be moved or split and the spatial index would not need to be rebuilt.

Spatial Cartridge release 8.03 introduced two distinct algorithms for building a spatial index: fixed-size tiling and variable-sized tiling. Based on testing and customer feedback, for release 8.0.4, Oracle recommends using only fixed-size tiling on production systems. Variable-sized tiling, while it has theoretical advantages in some situations, is included for experimentation purposes only.

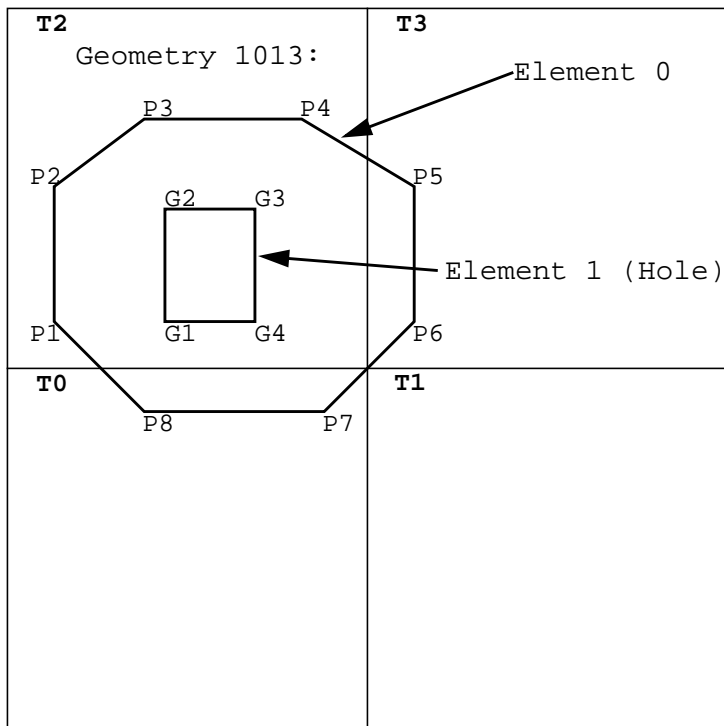
In spatial indexing, the object space (the layer where all geometric objects are located,) is subjected to a process called **tessellation**, which defines exclusive and exhaustive cover tiles of every stored element. Spatial Cartridge can use either fixed-size or variable-sized tiles to cover a geometry.

The number of tiles used to cover an element is a user-tunable parameter. Using either smaller fixed-size tiles or more variable-sized tiles provides a better fit of the tiles to the element. The fewer the number of tiles or the larger the tiles, the coarser the fit.

### 1.5.1 Tessellation of a Layer

The process of determining which tiles cover a given element is called tessellation. The tessellation process is a quad-tree decomposition, where the object space is broken down into four equal-sized covering tiles. Successive tessellations break those tiles down into four smaller tiles, and this process continues until the desired level has been achieved. The results of the tessellation process on an element are stored in the <layername>\_SDOINDEX table. See Section 2.3, “Index Creation” for more information on tessellation.

Figure 1–3 illustrates geometry 1013 tessellated to a maximum of four cover tiles. The cover tiles are then shown stored in the <layername>\_SDOINDEX table.

**Figure 1-3 Tessellated Figure**

Only three of the four tiles generated by the first tessellation interact with the geometry. Only those tiles that interact with the geometry are stored in the <layername>\_SDOINDEX table, as shown in Table 1-5. In this example, three fixed-size tiles are used.

**Table 1-5 <layername>\_SDOINDEX Using Fixed-Size Tiles**

SDO_GID <number>	SDO_CODE <raw>
1013	T0
1013	T2
1013	T3

All elements in a geometry are tessellated. In a multi-element polygon like 1013, Element 1 is already covered by tile T2 from the tessellation of Element 0.

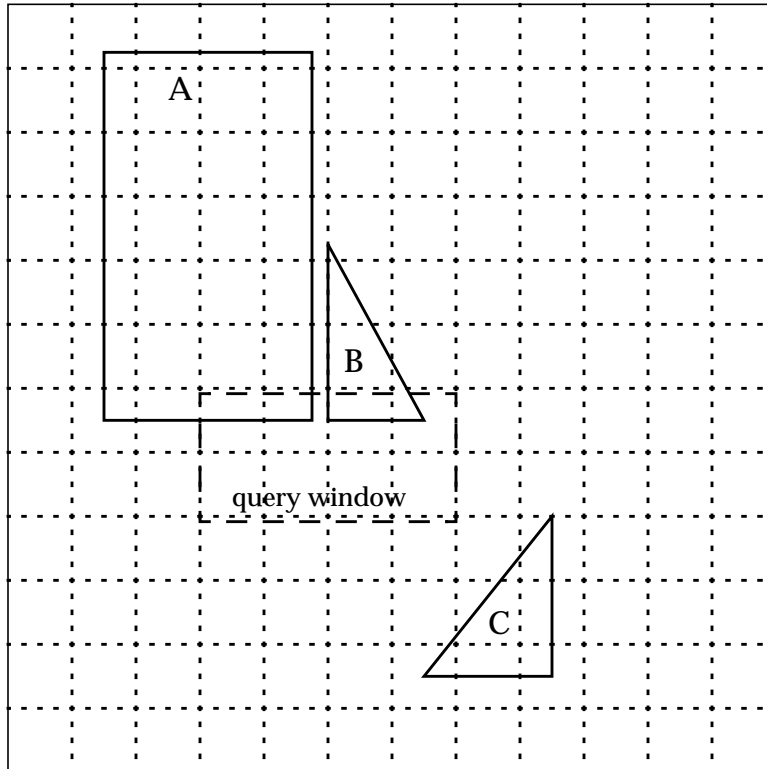
## 1.5.2 Fixed-Size Tile Spatial Indexing

Fixed-size tile spatial indexing is the recommended indexing method. This method uses cover tiles of equal size to cover a geometry. Because all the tiles are the same size, the standard SQL equality operator (=) can be used to compare tiles during a join operation. This results in excellent performance characteristics.

If you select a small fixed-size tile to cover small geometries and then try to use the same size tile to cover a very large geometry, a large number of tiles would be required, thereby increasing the size of the index table. However, if the fixed-size tile size chosen is large, so that fewer tiles are generated in the case of a large geometry, then the index selectivity suffers because the large tiles do not fit the small geometries very well. Figure 1-4 and Figure 1-5 illustrate the relationships between tile size, selectivity, and the number of cover tiles.

Using a small fixed-size tile as shown in Figure 1-4, selectivity is good, but a large number of tiles is needed to cover large geometries. A window query would easily identify geometries A and B, but would reject C.

**Figure 1-4 Fixed-Size Tiling with Many Small Tiles**

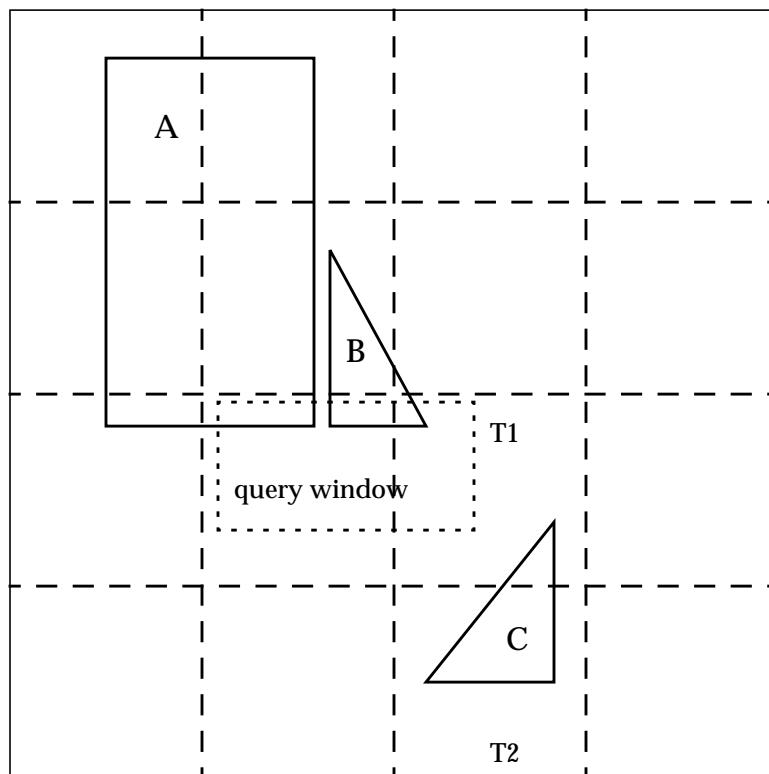




Using a large fixed-size tile as shown in Figure 1-5, fewer tiles are needed to cover the geometries, but the selectivity is poor. A window query would likely pick up all three geometries. Any object that shares tile T1 or T2 would identify object C as a candidate, even though the objects may be far apart, such as objects B and C are in this figure.

Use the `SDO_TUNE. ESTIMATE_TILING_LEVEL()` function to determine an appropriate tiling level for your data set.

**Figure 1-5 Fixed-Size Tiling with Fewer Large Tiles**



### 1.5.3 Variable-Sized Tile Spatial Indexing

---

---

**Note:** Variable-sized tile spatial indexing is not recommended for production environments. It is included primarily for experimentation purposes.


---

---

Variable-sized tile spatial indexing uses tiles of different sizes to approximate a geometry. The user specifies the number of tiles per object that should be used to approximate it and this governs the tiling process. As in the case of a linear quad tree, the cover tiles depend on the size and shape of each geometry being indexed and therefore good primary filter selectivity can be achieved. Figure 1–6 illustrates the approximation that variable-sized tiles can achieve.

In Figure 1-6, the variable-sized cover tiles conform closely to each geometry, resulting in good selectivity. The number of tiles needed to cover a geometry is controlled using the SDO\_NUMTILES column in the <layername>\_SDOLAYER table. See Section 2.3.3 for information on selecting appropriate values for variable-sized tiling.

**Figure 1-6 Variable-Sized Tile Spatial Indexing**



Two geometries may interact if a tile of one object is equal to, inside of, or contains a tile of the other. Thus, the query predicate to compare tiles involves a test for either equality or containment. This is unlike fixed-size tiling, which only requires an equality check. Example 1–1 demonstrates this feature (“5” is an arbitrary window identifier).

**Example 1–1**

```
SELECT r.sdo_gid
FROM roads_sdoindex r,
     window_sdoindex w
WHERE w.sdo_gid = 5
     AND (r.sdo_code BETWEEN w.sdo_code AND w.sdo_maxcode OR
          w.sdo_code BETWEEN r.sdo_code AND r.sdo_maxcode);
```


To reduce the number of times a complex predicate needs to be applied, variable-sized tile indexing uses a mechanism similar to spatial partitioning. To use this mechanism, select a tiling level, called the **groupcode level**, that results in tiles larger than any variable-sized tile generated for all the geometries in the layer or data set of interest. Each tile at the specified groupcode level can be considered a spatial partition. This reduces the size of the data set on which the complex predicate is evaluated. Example 1–2 illustrates this feature.

**Example 1–2**

```
SELECT r.sdo_gid
FROM layer_sdoindex r,
     window_sdoindex w
WHERE w.sdo_gid = 5
     AND r.sdo_group_code = w.sdo_groupcode
     AND (r.sdo_code BETWEEN w.sdo_code AND w.sdo_maxcode OR
          w.sdo_code BETWEEN r.sdo_code AND r.sdo_maxcode);
```

In Figure 1–7, consider the domain partitioned into 16 subregions. If a join compares tiles from the two objects, under normal circumstances the join operation would process tiles from the entire domain, searching for tiles that interact. However, if you constrain the processing to common partitions, then only partitions 5 and 6 would need to be processed. This may result in substantial performance improvements.

Figure 1–7 Spatially Partitioning Data



## 1.6 Partitioned Point Data

Spatial Cartridge has an enhanced spatial indexing mechanism capable of handling very large datasets consisting of complex geometries. For applications handling point datasets that are several tens of gigabytes or larger, further performance gains can be achieved by using Oracle8 table partitioning features.

Table partitioning is only available with the Partitioning Option of Oracle8 Enterprise Edition. If the Partitioning Option is available to you, the preferred method is to use Oracle8 table partitioning in conjunction with Spatial Cartridge spatial indexing (see the *Oracle8 Concepts* guide for a description of Oracle8 Partitioning). A technical white paper titled, “*Leveraging Oracle8 Partitioning and the Spatial Cartridge for*

*Large Point Datasets*,” describing the use of partitioning and spatial indexing for point datasets may be obtained from the Oracle corporate web site at:

<http://www.oracle.com/st/cartridges/spatial/collateral>

A previous release of Spatial Data Option (from which Spatial Cartridge has evolved) utilized its own version of table partitioning instead of spatial indexing. Chapter 4 briefly describes the old partitioning scheme for those customers with legacy point datasets. Any references to point data partitioning in this manual (such as the “Partitioned Point Data Procedures” section in Chapter 5) refer to this legacy feature. While this feature is still available in Spatial Cartridge, the preferred approach is to use Oracle8 Partitioning Option and spatial indexing.

---

# Loading Spatial Data

This chapter describes how to load spatial data into a database, including storing the data in a table and creating a spatial index for it.

## 2.1 Load Model

There are two steps involved in loading raw data into a spatial database such that it can be queried efficiently:

1. loading the data into spatial tables
2. creating or updating the index on the spatial tables

Table 2-1 through Table 2-4 show the format of the tables needed to store and index spatial data.

**Table 2-1** *<layername>\_SDOLAYER*

SDO_ORDCNT	SDO_LEVEL	SDO_NUMTILES	SDO_COORDSYS
<number>	<number>	<number>	<varchar>

**Table 2-2** *<layername>\_SDODIM Table or View*

SDO_DIMNUM	SDO_LB	SDO_UB	SDO_TOLERANCE	SDO_DIMNAME
<number>	<number>	<number>	<number>	<varchar>

**Table 2–3** *<layername>\_SDOGEOM Table or View*

SDO_GID	SDO_ESEQ	SDO_ETYPE	SDO_SEQ	SDO_X1	SDO_Y1	...	SDO_Xn	SDO_Yn
<number>	<number>	<number>	<number>	<number>	<number>	...	<number>	<number>

**Table 2–4** *<layername>\_SDOINDEX Table*

SDO_GID	SDO_CODE	SDO_MAXCODE	SDO_GROUPCODE	SDO_META
<number>	<raw>	<raw>	<raw>	<raw>

## 2.2 Load Process

The process of loading data can be classified into two categories:

- *bulk loading of data*  
This process is used to load large volumes of data into the database and uses the SQL\*Loader to load the data.
- *transactional inserts*  
This process is used to insert relatively small amounts of data into the database and is analogous to the INSERT statement in SQL.

### 2.2.1 Bulk Loading

Bulk loading can be used to import large amounts of legacy or ASCII data into a spatial database. Bulk loading is accomplished using the SQL\*Loader<sup>1</sup>.

Example 2–1 shows the format of the raw data and control file that would be required to load the data into the SDOGEOM table with the layer name ROADS. You can choose any format of ASCII data as long you can write a SQL\*Loader control file to load that data into the tables.

Assume that the ASCII data consists of a file with delimited columns, and separate rows fixed by the limits of the table with the following format.

**Example 2–1**

```
geometry rows:      GID, ESEQ, ETYPE, SEQ, LON1, LAT1, LON2, LAT2
```

<sup>1</sup> See the *Oracle Server Utilities User's Guide* for information on the SQL\*Loader.



The coordinates in the geometry rows represent the end points of line segments, which taken together, represent a polygon. Example 2-2 shows the control file for loading the data into the geometry table.

**Example 2-2**

```
LOAD DATA INFILE *
INTO TABLE ROADS_SDOGEOM
FIELDS TERMINATED BY WHITESPACE TRAILING NULLCOLS
(SDO_GID INTEGER EXTERNAL,
SDO_ESEQ INTEGER EXTERNAL,
SDO_ETYPE INTEGER EXTERNAL,
SDO_SEQ INTEGER EXTERNAL,
SDO_X1 FLOAT EXTERNAL,
SDO_Y1 FLOAT EXTERNAL,
SDO_X2 FLOAT EXTERNAL,
SDO_Y2 FLOAT EXTERNAL)

BEGINDATA
1 0 3 0 -122.401200 37.805200 -122.401900 37.805200
1 0 3 1 -122.401900 37.805200 -122.402400 37.805500
1 0 3 2 -122.402400 37.805500 -122.403100 37.806000
1 0 3 3 -122.403100 37.806000 -122.404400 37.806800
1 0 3 4 -122.404400 37.806800 -122.401200 37.805200
1 1 3 0 -122.405900 37.806600 -122.407549 37.806394
1 1 3 1 -122.407549 37.806394 -122.408300 37.806300
1 1 3 2 -122.408300 37.806300 -122.409100 37.806200
1 1 3 3 -122.409100 37.806200 -122.405900 37.806600
2 0 2 0 -122.410800 37.806000 -122.412300 37.805800
2 0 2 1 -122.412300 37.805800 -122.414100 37.805600
2 0 2 2 -122.414100 37.805600 -122.412300 37.805800
2 0 2 3 -122.412300 37.805800 -122.410800 37.806000
3 0 1 0 -122.567474 38.643564
3 0 1 1 -126.345345 39.345345
```

Note that table ROADS\_SDOGEOM exists in the schema before attempting the load.

In Example 2-3, the data resides in a single flat file and the data set consists of point, line string, and polygon data. The data uses fixed-position columns and over-loaded table rows.

**Example 2-3**

```
SDO_GID SDO_ESEQ SDO_ETYPE SDO_SEQ SDO_X1 SDO_Y1 SDO_X2 SDO_Y2
```

The corresponding control file for this format of input data would be:

```
LOAD DATA INFILE *
INTO TABLE NEW_SDOGEOM
(SDO_GID POSITION (1:5) INTEGER EXTERNAL,
SDO_ESEQ POSITION (7:10) INTEGER EXTERNAL,
SDO_ETYPE POSITION (12:15) INTEGER EXTERNAL,
SDO_SEQ POSITION (17:21) INTEGER EXTERNAL,
SDO_X1 POSITION (23:35) FLOAT EXTERNAL,
SDO_Y1 POSITION (37:48) FLOAT EXTERNAL,
SDO_X2 POSITION (50:62) FLOAT EXTERNAL,
SDO_Y2 POSITION (64:75) FLOAT EXTERNAL)

BEGINDATA
1 0 3 0 -122.401200 37.805200 -122.401900 37.805200
1 0 3 1 -122.401900 37.805200 -122.402400 37.805500
1 0 3 2 -122.402400 37.805500 -122.403100 37.806000
1 0 3 3 -122.403100 37.806000 -122.404400 37.806800
1 0 3 4 -122.404400 37.806800 -122.401200 37.805200
1 1 3 0 -122.405900 37.806600 -122.407549 37.806394
1 1 3 1 -122.407549 37.806394 -122.408300 37.806300
1 1 3 2 -122.408300 37.806300 -122.409100 37.806200
1 1 3 3 -122.409100 37.806200 -122.405900 37.806600
2 0 2 0 -122.410800 37.806000 -122.412300 37.805800
2 0 2 1 -122.412300 37.805800 -122.414100 37.805600
2 0 2 2 -122.414100 37.805600 -122.412300 37.805800
2 0 2 3 -122.412300 37.805800 -122.410800 37.806000
3 0 1 0 -122.567474 38.643564
3 0 1 1 -126.345345 39.345345
```

## 2.2.2 Transactional Insert Using SQL

Spatial Cartridge uses standard Oracle8 tables that can be accessed or loaded with standard SQL syntax. Example 2-4 loads data for a geometry (GID 17) consisting of a polygon with four sides that contains both a hole and point. Notice that the first coordinate of the polygon (5, 20) is repeated at the end to close the polygon.

### Example 2-4

```
INSERT INTO SAMPLE_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3,
SDO_Y3, SDO_X4, SDO_Y4, SDO_X5, SDO_Y5)
VALUES (17, 0, 3, 0, 5, 20, 5, 30, 10, 30, 10, 20, 5, 20);
```

```

-- hole
INSERT INTO SAMPLE_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
                             SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3,
                             SDO_Y3, SDO_X4, SDO_Y4, SDO_X5, SDO_Y5)
VALUES (17, 1, 3, 0, 8, 21, 8, 24, 9, 24, 9, 21, 8, 21);

-- point
INSERT INTO SAMPLE_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
                             SDO_X1, SDO_Y1)
VALUES (17, 2, 1, 0, 9, 29);

```

The SQL INSERT statement inserts one row of data per call. In Example 2-4, the table had enough columns to store the polygon in a single row. However, if your table had fewer columns (or your polygon had more points), you would have to perform multiple inserts in order to match the table structure; the data would not wrap automatically to the next row. To load a large geometry, repeat the SDO\_GID, SDO\_ESEQ, and SDO\_ETYPE, and increment the SDO\_SEQ for each line as shown in Example 2-5.

#### **Example 2-5**

```

INSERT INTO SAMPLE2_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
                             SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3,
                             SDO_Y3, SDO_X4, SDO_Y4, SDO_X5, SDO_Y5)
VALUES (18, 0, 3, 0, 1, 15, 1, 16, 2, 17, 3, 17, 4, 18);

INSERT INTO SAMPLE2_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
                             SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3,
                             SDO_Y3, SDO_X4, SDO_Y4, SDO_X5, SDO_Y5)
VALUES (18, 0, 3, 1, 4, 18, 5, 18, 6, 19, 7, 18, 6, 17);

INSERT INTO SAMPLE2_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
                             SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3,
                             SDO_Y3, SDO_X4, SDO_Y4, SDO_X5, SDO_Y5)
VALUES (18, 0, 3, 2, 6, 17, 7, 16, 7, 15, 6, 14, 7, 13);

INSERT INTO SAMPLE2_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
                             SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3,
                             SDO_Y3, SDO_X4, SDO_Y4, SDO_X5, SDO_Y5)
VALUES (18, 0, 3, 3, 7, 13, 6, 12, 5, 13, 4, 13, 3, 14);

INSERT INTO SAMPLE2_SDOGEOM (SDO_GID, SDO_ESEQ, SDO_ETYPE, SDO_SEQ,
                             SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3,
                             SDO_Y3)

```

```
VALUES (18, 0, 3, 4, 3, 14, 2, 14, 1, 15);
```

### 2.2.3 Transactional Insert Using Spatial Geometry Functions

Spatial Cartridge provides two functions to facilitate inserting data into spatial tables. A benefit to using these functions is that the issue of row-wrapping when loading elements with multiple points is handled automatically by these functions.

There are two steps to incrementally add data to the spatial tables:

1. Initialize the element that needs to be stored. Note that this process does not fill in any coordinate information for the element. Two parameters are passed to the `SDO_GEOM.INIT_ELEMENT()` function, which initializes the element:

- the name of the layer (for example, `ROADS`)
- a `GID` that is a unique identifier for the geometry

The `SDO_GEOM.INIT_ELEMENT()` function returns the sequence number of the element in the geometry. This sequence number is required as a parameter to the `SDO_GEOM.ADD_NODES()` procedure.

2. Fill in the coordinate information for the element using the `SDO_GEOM.ADD_NODES()` procedure. This procedure takes the following parameters:

- the name of the layer
- the `GID`
- the sequence number of the element
- the element type
- the list of vertices in the geometry, specified as a series of `X,Y` coordinate pairs

Note that you must explicitly close a polygon by repeating the coordinates of the first vertex as the last vertex.

In Example 2-6, a simple polygon, geometry number 1234, consisting of five vertices needs to be stored. The first step is to call `SDO_GEOM.INIT_ELEMENT()` to initialize the element.

#### **Example 2-6**

```
elem_value := sdo_geom.init_element('ROADS', 1234);
```

Next, call `SDO_GEOM.ADD_NODES()` to fill in the attributes of the polygon. The vertices can be added in either clockwise or counter-clockwise order.


```
sdo_geom.add_nodes('ROADS', 1234, elem_value, sdo_geom.polygon_type, Ax,
Ay, Bx, By, Cx, Cy, Dx, Dy, Ex, Ey, Ax, Ay);
```

Close the polygon by repeating the first vertex ( $A_x, A_y$ ) as the last vertex.

In Example 2-7, assume that the geometry shown in Figure 2-1 needs to be stored. The geometry consists of a polygon with a hole in it. Note that both calls to the `SDO_GEOM.ADD_NODES()` procedure are made with the same `GID (6789)` because this is a single object even though it is composed of two elements.

**Figure 2-1 Polygon with a Hole**

Geometry 6789:



**Example 2-7**

```
val1 := sdo_geom.init_element('PARKS', 6789);
sdo_geom.add_nodes('PARKS', 6789, val1, SDO_GEOM.POLYGON_TYPE, P1_x, P1_y,
P2_x, P2_y, P3_x, P3_y, P4_x, P4_y, P5_x, P5_y, P6_x, P6_y, P1_x, P1_y);
val2 := sdo_geom.init_element('PARKS', 6789);
sdo_geom.add_nodes('PARKS', 6789, val2, SDO_GEOM.POLYGON_TYPE, G1_x, G1_y,
G2_x, G2_y, G3_x, G3_y, G4_x, G4_y, G1_x, G1_y);
```

## 2.3 Index Creation

Once data has been loaded into the spatial tables through either bulk or transactional loading, a spatial index needs to be created on the tables for efficient access to the data.

Create an Oracle8 table called `<layername>_SDOINDEX` as follows:

```
SQL> create table <layername>_SDOINDEX
2  (
3    SDO_GID integer,
4    SDO_CODE raw(255)
5  );
```

For a bulk load, you can call the `SDO_ADMIN.POPULATE_INDEX()` procedure once to tessellate the geometry table and add the generated tiles to the spatial index table. The argument to this procedure is simply the name of the layer. The level to which the geometry should be tessellated, and whether to use the fixed-size or variable-sized tile indexing technique is determined by values in the `<layername>_SDOLAYER` table.

If data is updated in or deleted from a specific geometry table, you can call `SDO_ADMIN.UPDATE_INDEX()` to update the index for one `SDO_GID`. The arguments to this procedure are the name of the layer and the `SDO_GID` of the designated geometry.

See Chapter 5, “Administrative Procedures” for a complete description of the `SDO_ADMIN.POPULATE_INDEX()` and `SDO_ADMIN.UPDATE_INDEX()` procedures.

### 2.3.1 Choosing a Tessellation Algorithm

Spatial Cartridge provides two methods for spatial indexing. Fixed-size tiling is recommended for all production applications. For advanced development applications, you may want to experiment with variable-sized tiling, which theoretically could provide better selectivity in some data sets.

Which tessellation algorithm is used by the `SDO_ADMIN.POPULATE_INDEX()` and `SDO_ADMIN.UPDATE_INDEX()` procedures is determined by the values of the `SDO_LAYER` and `SDO_NUMTILES` columns in the `<layername>_SDOLAYER` table as follows:

<code>SDO_LEVEL</code>	<code>SDO_NUMTILES</code>	Action
NULL	NULL	Error

SDO_LEVEL	SDO_NUMTILES	Action
>= 1	NULL	Fixed-size tiling
>= 1	>= 1	Indexing with variable-sized tiles. The SDO_LEVEL column defines the partition bucket size. The SDO_NUMTILES column defines the number of tiles to generate per geometry. Note: variable-sized tiling is for experimentation purposes only.
NULL	>= 1	Not supported


### 2.3.2 Spatial Indexing with Fixed-Size Tiles

Oracle recommends using fixed-size cover tiles for indexing a geometry.

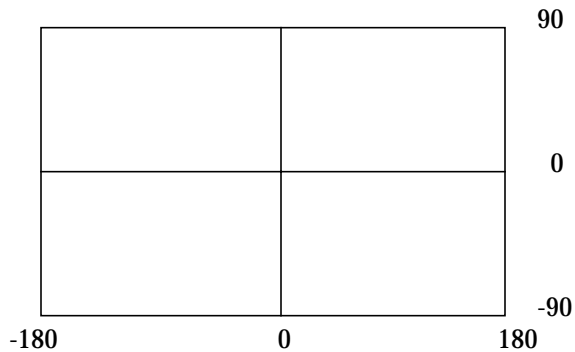
The fixed-size tile algorithm is expressed as a level referring to the number of tessellations performed. To use fixed-size tile indexing, set the SDO\_NUMTILES column in the <layername>\_SDOLAYER table to NULL and the SDO\_LEVEL column to the desired tiling level. The relationship between the tiling level and the resulting size of the tiles is dependent on the domain of the layer.

The domain used for indexing is defined by the upper and lower boundaries of each dimension stored in the <layername>\_SDODIM table. A typical domain in a GIS application could be -90 to 90 degrees for latitude, and -180 to 180 degrees for longitude<sup>1</sup>, as represented in Figure 2-2.

<sup>1</sup>The transference of the domain onto a sphere or Mercator projection is left to GIS (or other) application programmers. Spatial Cartridge treats the domain as a conventional X by Y rectangle.


**Figure 2-2 Sample GIS Domain**

If the `SDO_LEVEL` column is set to 1, then the tiles created by the indexing mechanism are the same size as tiles at the first level of tessellation. Each tile would be 180 degrees by 90 degrees as shown in Figure 2-3.

**Figure 2-3 Fixed-Size Tiling at Level 1**

The formula for the number of fixed-size tiles is  $4^n$  where  $n$  is the number of tessellations, stored in the `SDO_LEVEL` column. Figure 2-4 shows fixed-size tiling at level 2. In this figure, each tile is 90 degrees by 45 degrees.



**Figure 2–4 Fixed-Size Tiling at Level 2**

The size of a tile can be determined by applying the following formula to each dimension:

$$\text{length} = (\text{upper\_bound} - \text{lower\_bound}) / 2^{\text{sdo\_level}}$$

The length refers to the length of the tile along the specified dimension. Applying this formula to the tiling shown in Figure 2–4 yields the following sizes:

$$\begin{aligned} \text{length for dimension X} &= (180 - (-180)) / 2^2 \\ &= (360) / 4 \\ &= 90 \end{aligned}$$

$$\begin{aligned} \text{length for dimension Y} &= (90 - (-90)) / 2^2 \\ &= (180) / 4 \\ &= 45 \end{aligned}$$

Thus, at level 2 the tiles are 90x45 degrees in size. As the number of levels increases, the tiles become smaller and smaller. Smaller tiles provide a more precise fit of the tiles over the geometry being indexed. However, because the number of tiles generated is unbounded, you must take into account the performance implications of using higher levels. The `SDO_TUNE.ESTIMATE_TILING_LEVEL()` function can be used to determine an appropriate level for indexing with fixed-size tiles. See Chapter 6 for a description of this procedure.

Besides the performance aspects related to selecting a fixed-size tile, tessellating the geometry into fixed-size tiles might have benefits related to the type of data being stored, such as using tiles sized to represent 1-acre farm plots, city blocks, or individual pixels on a display. Data modeling is an important part any database design, and is essential in a spatial database where the data often represents actual physical locations.

**Example 2–8**

Assume that data has been loaded into a layer called ROADS. To create a spatial index, create a table ROADS\_SDOINDEX and invoke the following procedure:

```
sdo_admin.populate_index('ROADS');
```

The value in the SDO\_LEVEL column of ROADS\_SDOLAYER can be used as a tuning parameter while tessellating objects. Increasing the level increases the number of tiles to provide a more precise fit of the tiles over the object. See the description of the ESTIMATE\_TILING\_LEVEL( ) function in Chapter 6 for information on estimating the tiling level in several different ways.

After SDO\_ADMIN.POPULATE\_INDEX( ) has been called to fill the spatial index, you should also create standard indexes on the <layername>\_SDOINDEX.SDO\_CODE column.

If a geometry with an SDO\_GID 5944 has been added to the spatial tables, update the index with the following procedure:

```
sdo_admin.update_index('ROADS', 5944);
```

SDO\_ADMIN.POPULATE\_INDEX( ) and SDO\_ADMIN.UPDATE\_INDEX( ) behave differently than the CREATE INDEX statement in SQL. An implicit commit is not executed after the procedures are called. Therefore these transactions can be rolled back.

SDO\_ADMIN.POPULATE\_INDEX operates as a single transaction. To reduce the amount of rollback required to execute this procedure, you can write a routine that loops and calls SDO\_ADMIN.UPDATE\_INDEX( ). See Section A.1.1.1, “cr\_spatial\_index.sql Script” for more information.

### 2.3.3 Spatial Indexing with Variable-Sized Tiles

Spatial indexing with variable-sized tiles is not recommended for production systems. Variable-sized tiling is included in Spatial Cartridge primarily for experimentation purposes.

To use variable-sized tiling, the SDO\_LEVEL and SDO\_NUMTILES columns must be set in the <layername>\_SDOLAYER table.

The SDO\_NUMTILES column determines the number of tiles that will be used to cover a geometry being indexed. Typically this value is small, such as 4 or 8 tiles. However, the larger the number of tiles, the better the tiles will fit the geometry

being covered. This increases the selectivity of the primary filter. See Section 3.3.2 and Section 3.3.3 for a discussion of primary and secondary filters.

The `SDO_LEVEL` column indicates the spatial partitioning level for the generated tiles. See Section 1.5.3 for a description of the spatial partitioning utilized by Spatial Cartridge when using variable-sized tiles.

Setting the proper `SDO_LEVEL` value is more art than science. One approach would be use the `SDO_TUNE.ESTIMATE_TILING_LEVEL()` function to determine an appropriate starting `SDO_LEVEL` value, and then compare the performance with slightly higher or lower values.



---

# Querying Spatial Data

This chapter describes how the structures of a Spatial Cartridge layer are used to resolve spatial queries and spatial joins. For the sake of clarity, the examples all use fixed-size tiling.

## 3.1 Query Model

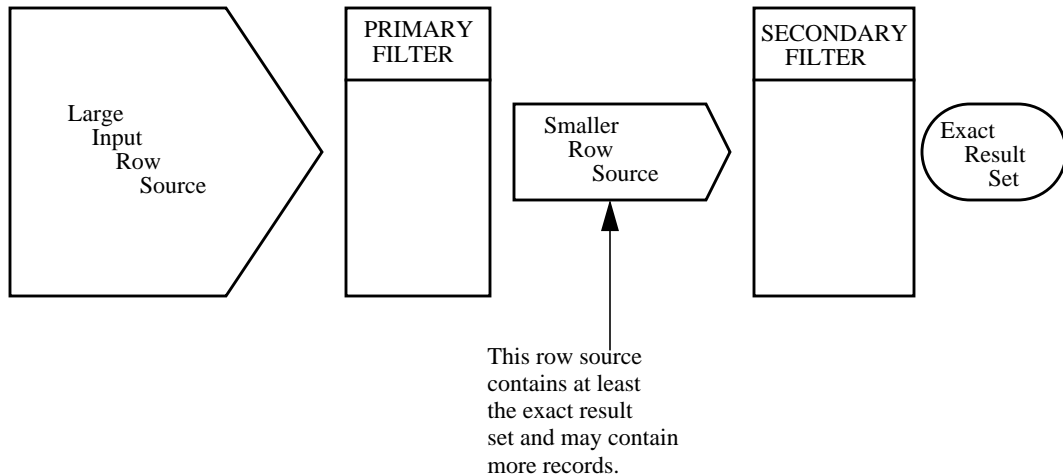
Spatial Cartridge uses a *two-tier* query model to resolve spatial queries and spatial joins. The term is used to indicate that two distinct operations are performed in order to resolve queries. The output of both operations yields the exact result set.

The two operations are referred to as *primary* and *secondary* filter operations.

- The primary filter permits fast selection of a small number of candidate records to pass along to the secondary filter. The primary filter uses approximations in order to reduce computational complexity and is considered a lower cost filter.
- The secondary filter applies exact computational geometry to the result set of the primary filter. These exact computations yield the final answer to a query. The secondary filter operations are computationally more expensive, but they are only applied to the relatively small result set from the primary filter.

Figure 3–1 illustrates the relationship between the primary and secondary filters.

Figure 3–1 Query Model



Spatial Cartridge uses a spatial index to implement the primary filter. This is described in detail in following sections.

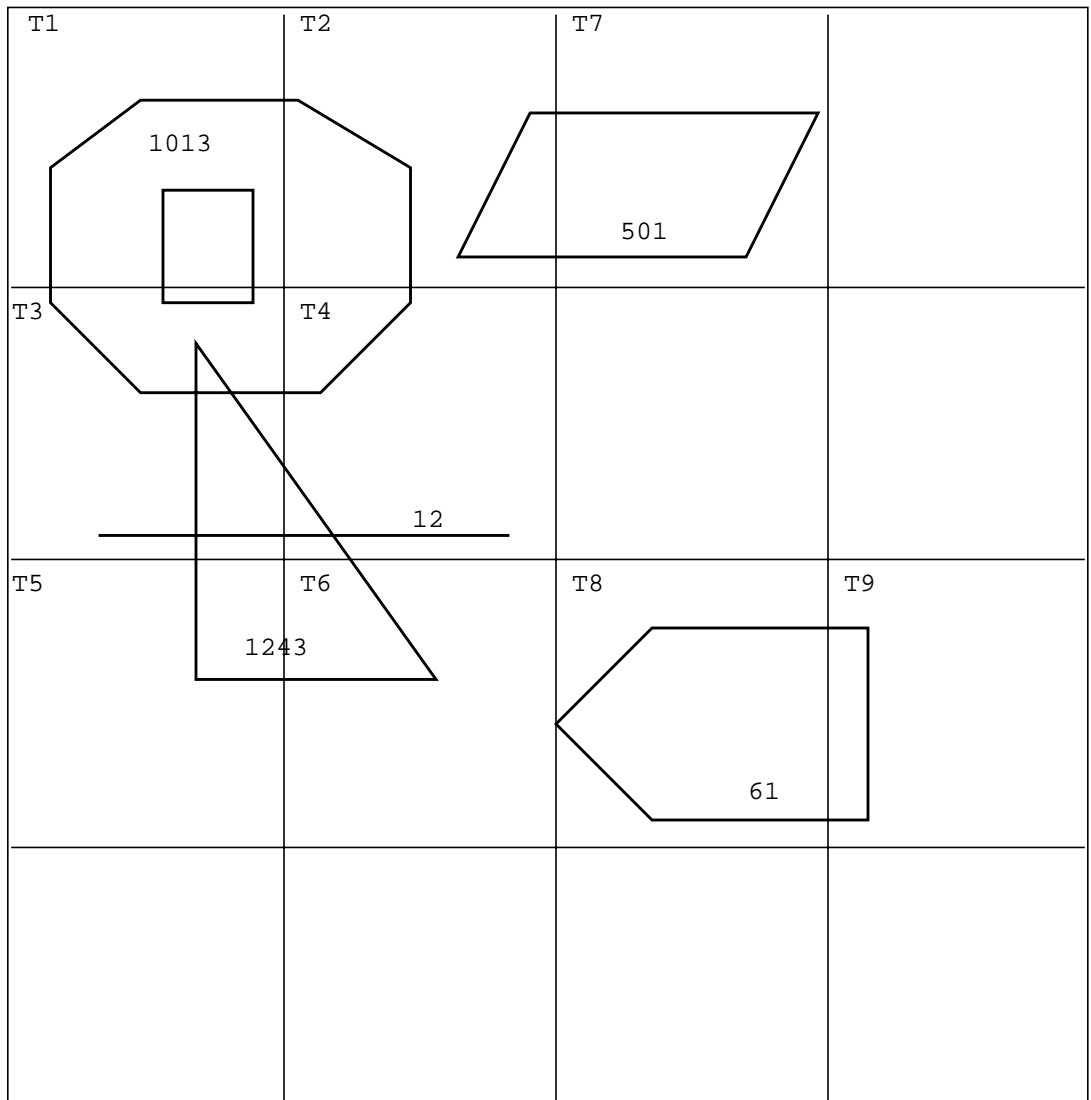
A function used as a secondary filter is `SDO_GEOM.RELATE( )`, which determines the spatial relationship between two given geometries, such as whether they touch, overlap, or if one is inside the other.

Spatial Cartridge does not require the use of both the primary and secondary filters. In some cases, just using the primary filter is sufficient. For example, a *zoom* feature in a mapping application queries for data that overlaps a rectangle representing visible boundaries. The primary filter very quickly returns a superset of the query. The mapping application can then apply clipping routines to display the target area.

## 3.2 Spatial Data Model

An important concept in the spatial data model is that each element is represented in the `<layername>_SDOINDEX` table by a set of exclusive and exhaustive tiles. This means that the tiles fully cover the object (**exhaustive**) and that no tiles overlap each other (**exclusive**).

Consider the following layer containing several objects in Figure 3–2. Each object is labeled with its `SDO_GID`. The relevant tiles are labeled with 'Tn'.

**Figure 3-2 Tessellated Layer with Multiple Objects**

The Spatial Cartridge layer tables would have the following information stored in them for these geometries as shown in Table 3-1, Table 3-2, and Table 3-3.

Table 3-1 &lt;layername&gt;\_SDOLAYER

SDO_ORDCNT (number)	SDO_LEVEL (number)	SDO_NUMTILES (number)
4	2	NULL

Table 3-2 &lt;layername&gt;\_SDOGEOM

SDO_GID (number)	SDO_ESEQ (number)	SDO_ETYPE (number)	SDO_SEQ (number)	SDO_X1 (number)	SDO_Y1 (number)	SDO_X2 (number)	SDO_Y2 (number)
1013	0	3	0	P1(X)	P1(Y)	P2(X)	P2(Y)
1013	0	3	1	P2(X)	P2(Y)	P3(X)	P3(Y)
1013	0	3	2	P3(X)	P3(Y)	P4(X)	P4(Y)
1013	0	3	3	P4(X)	P4(Y)	P5(X)	P5(Y)
1013	0	3	4	P5(X)	P5(Y)	P6(X)	P6(Y)
1013	0	3	5	P6(X)	P6(Y)	P7(X)	P7(Y)
1013	0	3	6	P7(X)	P7(Y)	P8(X)	P8(Y)
1013	0	3	7	P8(X)	P8(Y)	P1(X)	P1(Y)
1013	1	3	0	G1(X)	G1(Y)	G2(X)	G2(Y)
1013	1	3	1	G2(X)	G2(Y)	G3(X)	G3(Y)
1013	1	3	2	G3(X)	G3(Y)	G4(X)	G4(Y)
1013	1	3	3	G4(X)	G4(Y)	G1(X)	G1(Y)
501	0	3	0	A1(X)	A1(Y)	A2(X)	A2(Y)
501	0	3	1	A2(X)	A2(Y)	A3(X)	A3(Y)
501	0	3	2	A3(X)	A3(Y)	A4(X)	A4(Y)
501	0	3	3	A4(X)	A4(Y)	A1(X)	A1(Y)
1243	0	3	0	B1(X)	B1(Y)	B2(X)	B2(Y)
1243	0	3	1	B2(X)	B2(Y)	B3(X)	B3(Y)
1243	0	3	2	B3(X)	B3(Y)	B1(X)	B1(Y)
12	0	2	0	D1(X)	D1(Y)	D2(X)	D2(Y)
61	0	3	0	C1(X)	C1(Y)	C2(X)	C2(Y)
61	0	3	1	C2(X)	C2(Y)	C3(X)	C3(Y)
61	0	3	2	C3(X)	C3(Y)	C4(X)	C4(Y)
61	0	3	3	C4(X)	C4(Y)	C5(X)	C5(Y)
61	0	3	4	C5(X)	C5(Y)	C1(X)	C1(Y)




**Table 3–3** <layername>\_SDOINDEX

SDO_GID (number)	SDO_CODE (raw)
1013	T1
1013	T2
1013	T3
1013	T4
501	T2
501	T7
1243	T3
1243	T4
1243	T5
1243	T6
12	T3
12	T4
61	T8
61	T9

### 3.3 Spatial Query

A typical spatial query is to request all objects that lie within a defined fence or window. A **query window** is shown in Figure 3–3 by the dotted line box. A dynamic query window refers to a fence that is not defined in the database, but that must be defined and indexed prior to using it.

**Figure 3-3 Tessellated Layer with a Query Window**



### 3.3.1 Dynamic Query Window

If a query window does not already exist in the database, you must first insert it and create an index for it. Because not all Oracle users necessarily have insert privi-

leges, Spatial Cartridge includes the SDO\_WINDOW PL\*SQL package. See Chapter 8, “Window Functions” for more information.

The SDO\_WINDOW package is not automatically installed when you install Spatial Cartridge. This allows a DBA to control the schema under which this package operates. Choose an Oracle user who has insert privilege and compile the SDO\_WINDOW package under that user. For example, you could choose the mdsys Oracle user:

```
sqlplus mdsys/password
SQL> @$ORACLE_HOME/md/admin/sdowin.sql
SQL> @$ORACLE_HOME/md/admin/prvtwin.plb
```

After compiling, the routines are available for use. When you call a routine in this package, and the routine performs an INSERT operation, the insert will occur under the mdsys schema. Note that it is not a requirement to use the mdsys account. You can select any Oracle user with insert privileges.

If you need to perform other INSERT, UPDATE, or DELETE operations, and you cannot guarantee that the user of your application has those privileges, you can write your own PL\*SQL package similar to the SDO\_WINDOW package. You will have to compile your package under a user with the required database privileges.

### 3.3.2 Primary Filter

To resolve the window query shown in Figure 3–3, build a layer for the query fence if it is not already defined:

```
SQL> EXECUTE MDSYS.SDO_WINDOW.CREATE_WINDOW_LAYER (fencelayer, DIMNUM1, LB1,
UB1, TOLERANCE1, DIMNAME1, DIMNUM2, LB2, UB2, TOLERANCE2, DIMNAME2);
```

Next, insert the ordinates for the query fence into the layer tables:

```
SQL> EXECUTE DBMS_OUTPUT.PUTLINE(TO_CHAR(MDSYS.SDO_WINDOW.BUILD_WINDOW_FIXED
(comp_user, fencelayer, SDO_ETYPE, TILE_SIZE, X1, Y1, X2, Y2, X3, Y3, X4, Y4,
X1, Y1)));
```

Query SDO\_LEVEL from the <fencelayer>\_SDOLAYER table to pass the correct TILE\_SIZE to the SDO\_WINDOW.BUILD\_WINDOW\_FIXED() procedure.

Now you can construct a query that joins the index of the query window to the appropriate layer index and determines all elements that have these tiles in common. The following SQL query form is used:

```

SELECT DISTINCT A.SDO_GID
FROM <layer1>_SDOINDEX A, <fencelayer>_SDOINDEX B
WHERE A.SDO_CODE = B.SDO_CODE
      AND B.SDO_GID = {GID returned from SDO_WINDOW.BUILD_WINDOW_FIXED};

```

The result set of this query is the primary filter set. In this case, the result set is:

```
{ 1013,501,1243,12 }
```

### 3.3.3 Secondary Filter

The secondary filter performs exact geometry calculations of the tiles selected by the primary filter. The following example shows the primary and secondary filters:

```

SELECT SDO_X1, SDO_Y1, SDO_X2, SDO_Y2, SDO_X3, SDO_Y3, SDO_X4, SDO_Y4
FROM <layer1>_SDOGEOM,
(
SELECT SDO_GID GID1
  FROM (
        SELECT DISTINCT A.SDO_GID
        FROM <layer1>_SDOINDEX A,
             <fencelayer>_SDOINDEX B
        WHERE A.SDO_CODE = B.SDO_CODE
              AND B.SDO_GID = {GID returned from SDO_WINDOW.BUILD_WINDOW_FIXED}
      )
  WHERE SDO_GEOM.INTERACT('<layer1>', GID1, '<fence>', 1) = 'TRUE'
)
WHERE SDO_GID = GID1;

```

This query would return all the geometry IDs that lie within or overlap the window. In this example, the results of the secondary filter would be:

```
{1243,1013}
```

The example in this section uses the `SDO_GEOM.INTERACT()` secondary filter. Both the `INTERACT()` and `RELATE()` secondary filters are overloaded functions. For better performance, use the versions that explicitly list the coordinates of the query window whenever possible. See Chapter 7, “Geometry Functions” for details on using these functions.

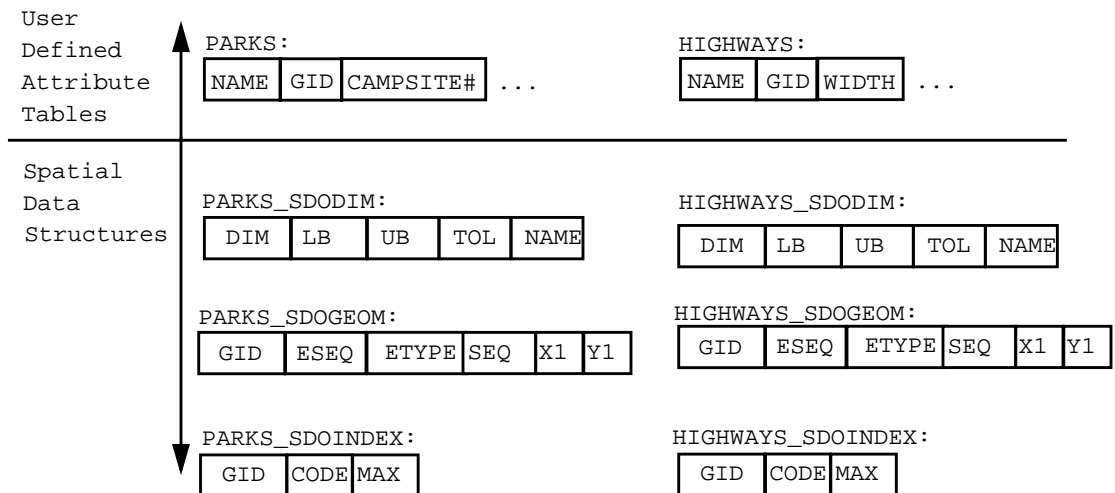
## 3.4 Spatial Join

A spatial join is the same as a regular join except that the predicate involves a spatial operator. In Spatial Cartridge, a spatial join takes place between two layers; specifically, two <layername>\_SDOINDEX tables are joined.

Spatial joins can be used to answer questions such as, “which highways cross national parks?”

This query could be resolved by joining a layer that stores national park geometries with one that stores highway geometries. Figure 3–4 illustrates how the join would be accomplished for this example.

**Figure 3–4 Spatial Join of Two Layers**



The PRIMARY filter would identify pairs of park GIDs and highway GIDs that cross in the index. The query that performs the primary filter join (assuming fixed-size tile indexing) is as follows:

```
SELECT DISTINCT A.SDO_GID,B.SDO_GID
FROM PARKS_SDOINDEX A, HIGHWAYS_SDOINDEX B
WHERE A.SDO_CODE = B.SDO_CODE
```

The result set of the primary filter must be passed through the secondary filter to get the exact set of parks/highways GID pairs that cross. The full query is shown in the following example:

```

SELECT DISTINCT SDO_GID
FROM (
  SELECT /*+ index(a PARKS_SDOINDEX_SDO_CODE_INDEX)
          index(b HIGHWAYS_SDOINDEX_SDO_CODE_INDEX)
          use_nl(a b)
          no_merge */
        DISTINCT A.SDO_GID GID_A, B.SDO_CODE GID_B
  FROM PARKS_SDOINDEX A, HIGHWAYS_SDOINDEX B
  WHERE A.SDO_CODE = B.SDO_CODE
)
WHERE SDO_GEOM.RELATE ('PARKS', GID_A,
                      'ANYINTERACT',
                      'HIGHWAYS', GID_B) <> 'FALSE';

```

The diagram illustrates the flow of data in the SQL query. A bracket on the right side of the subquery (the inner SELECT statement) is labeled "Primary Filter". An arrow points from this bracket to the subquery. Another bracket on the right side of the WHERE clause is labeled "Secondary Filter". An arrow points from this bracket to the WHERE clause. A third arrow points from the subquery to the WHERE clause, indicating that the results of the primary filter are passed through the secondary filter.

Suppose the original query had asked, “which *4-lane* highways cross national parks?” You could modify the preceding SQL statement to join back to the HIGHWAYS table where HIGHWAYS.WIDTH=4. This combination of spatial and relational attributes in a single query is one of the essential reasons for using Spatial Cartridge.

---

# Partitioning Point Data

Spatial Cartridge provides the essential functions, procedures, and scripts for using and managing both spatially indexed data and partitioned point data. The information in this chapter is relevant only to users utilizing table partitioning for very large quantities of point data.

## 4.1 Overview

Partitioning is a technique where data is loaded into tables that automatically subdivide when a predefined maximum size is reached. During subdivision, data is moved from the parent partition to the child partitions and the parent partition is dropped. Storage parameters for child partitions are inherited from the root partition and can be changed at any time.

A partitioned table has a partition key that is an HHCODE column created by encoding multidimensional point data using the `SDO_ENCODE()` function. In the partitioning process, at each subdivision, data is subdivided into  $2^n$  partitions where  $n$  is the number of dimensions encoded in the HHCODE column. You can encode up to 32 dimensions using Spatial Cartridge.

## 4.2 Partitioning Process

This manual does not attempt to provide the information necessary for fully utilizing table partitioning for point data. The following is a high-level description of the partitioning process:

1. Start with an Oracle8 table containing multidimensional point data. For example, columns of X and Y coordinate data from a blueprint or map.
2. Create a table or view from the original Oracle8 table containing the columns you want, plus a new HHCODE column.

An HHCODE column is a new data type used to encode multiple dimensions into a unique orderable value. HHCODE is not a point, but rather a bounded cell representing an object space in as many dimensions as have been defined. An HHCODE data type is defined as RAW(255).

3. Create the HHCODE data type by encoding multiple dimensions into a single value using the `SDO_ENCODE()` function. The HHCODE data type will be used as the partition key.
4. Register a partitioned table in the Spatial Cartridge data dictionary using the `SDO_ADMIN.REGISTER_PARTITION_INFO()` procedure. This procedure takes the name of a table, the name of the partition key column, and the maximum number of records you want stored in a partition before it subdivides.
5. Call the `SDO_ADMIN.PARTITION()` procedure with the name of the table or view containing the partition key column and the tablespace in which the partitions should be created. In this step, the data is partitioned based on dimensions encoded in the HHCODE column.
6. If the underlying table has constraints, grants, or triggers, the owner needs to use the `SDO_ADMIN.PROPAGATE_GRANTS()` procedure to set those properties on the partitions.
7. To add more partitioned point data, load the data into a table, and call `SDO_ADMIN.PARTITION()` again. The dimensions encoded in the HHCODE column must have the same boundaries to be loaded into the existing partitioned table.
8. After you have added data multiple times, or after adding or deleting a large amount of data, there may be partitions that exceed the high-water mark or there may be partitions that can be merged. Call the `SDO_ADMIN.REPARTITION()` procedure to reorganize the partitioned table. Repartitioning is a computation-intensive task that should be performed only when necessary.

## 4.3 Function Details

See the following sections for details of the functions supporting partitioned point data:

- Chapter 5, “Administrative Procedures”
  - `SDO_ADMIN.ALTER_HIGH_WATER_MARK`
  - `SDO_ADMIN.DROP_PARTITION_INFO`
  - `SDO_ADMIN.PARTITION`



- SDO\_ADMIN.PROPAGATE\_GRANTS
- SDO\_ADMIN.REGISTER\_PARTITION
- SDO\_ADMIN.REPARTITION
- SDO\_ADMIN.VERIFY\_PARTITIONS
- Chapter 9, “Partitioned Point Data Functions”
  - SDO\_BVALUETODIM
  - SDO\_COMPARE
  - SDO\_DATETODIM
  - SDO\_DECODE
  - SDO\_ENCODE
  - SDO\_TO\_BVALUE
  - SDO\_TO\_DATE
- Appendix A, “Sample SQL Scripts and Tuning Tips”
  - altpart.sql
  - drppart.sql
  - sdogrant.sql



---

## Administrative Procedures

The SDO\_ADMIN procedures create and maintain spatial structures in the database, and are used to perform the following tasks:

- tessellate entries in a geometry table and place them in a spatial index table
- register and manipulate partitioned spatial tables (partitioned tables are used only for large volumes of point data)
- verify spatial index and partitioned spatial table information

This chapter contains descriptions of the administrative procedures used for working with either spatially indexed geometric data or partitioned point data. These data structures are mutually exclusive and the procedures only work with the structure for which they are designed.

Table 5-1 lists the administrative procedures for working with spatially indexed geometry-based data. Table 5-2 later in this chapter lists procedures for working with partitioned point data.

**Table 5-1 Administrative Procedures for Spatially Indexed Data**

Procedure	Data Structure	Description
SDO_ADMIN.POPULATE_INDEX	Geometric objects	Generates a spatial index for the geometry table using either a set number of tiles, or a fixed-size tile.
SDO_ADMIN.SDO_CODE_SIZE	Geometric objects	Determines the required sizes for SDO_CODE and SDO_MAXCODE.
SDO_ADMIN.UPDATE_INDEX	Geometric objects	Updates the spatial index based on changes to the geometry table.
SDO_ADMIN.VERIFY_LAYER	Geometric objects	Checks for the existence of geometry and spatial index tables.

## SDO\_ADMIN.POPULATE\_INDEX

### Purpose

This procedure tessellates a list of geometric objects created by selecting all the entries in the geometry table that do not have corresponding entries in the spatial index table.

This procedure can generate either fixed or variable-sized tiles depending on values stored in the <layername>\_SDOLAYER table.

### Syntax

SDO\_ADMIN.POPULATE\_INDEX (*layername*)

### Keywords and Parameters

*layername* Specifies the name of the data set layer. The layername is used to construct the names of the geometry and spatial index tables.  
Data type is VARCHAR2.

### Usage Notes

Consider the following when using this procedure:

- The <layername>\_SDOINDEX table must be created prior to calling this procedure. Use the SQL CREATE TABLE statement to create the spatial index table.
- For performance reasons, create an index on the SDO\_GID column in the <layername>\_SDOGEOM table before calling this routine.
- This procedure generates either fixed-size or variable-sized tiles depending on values stored in the <layername>\_SDOLAYER table as follows:

SDO_LEVEL	SDO_NUMTILES	Action
NULL	NULL	Error
>= 1	NULL	Fixed-size tiling

SDO_LEVEL	SDO_NUMTILES	Action
>= 1	>= 1	Indexing with variable-sized tiles. The SDO_LEVEL column defines the partition bucket size. The SDO_NUMTILES column defines the number of tiles to generate per geometry.  Note: variable-sized tiling is for experimentation purposes only.
NULL	>= 1	Not supported

- If the <layername>\_SDOINDEX table is empty, the procedure selects all the geometries in the geometry table and generates index entries for them. If the index table is not empty, the procedure determines which entries in the geometry table do not have index entries, and generates them.
- SDO\_ADMIN.POPULATE\_INDEX( ) behaves similarly to the CREATE INDEX statement in SQL. An implicit commit is executed after the procedure is called.
- SDO\_ADMIN.POPULATE\_INDEX( ) operates as a single transaction. To reduce the amount of rollback required to execute this procedure, you can write a routine that loops and calls SDO\_ADMIN.UPDATE\_INDEX( ) repeatedly. See Section A.1.1.1, “cr\_spatial\_index.sql Script” for more information.

Example 5–1 tessellates all the geometric objects in the LAYER1\_SDOGEOM table and adds the generated tiles to the LAYER1\_SDOINDEX table.

**Example 5–1**

```
SQL> EXECUTE SDO_ADMIN.POPULATE_INDEX('layer1');
SQL> COMMIT;
```

## Related Topics

- SDO\_ADMIN.UPDATE\_INDEX() procedure

---

## SDO\_ADMIN.POPULATE\_INDEX\_FIXED

### Purpose

This procedure is provided for compatibility with Spatial Cartridge release 8.0.3 tables. This procedure has been replaced by enhanced features in the `SDO_ADMIN.POPULATE_INDEX()` procedure, and by supporting schema changes as shown in Section 1.4.

This procedure tessellates a list of geometric objects created by selecting all the entries in the geometry table that do not have corresponding entries in the spatial index table. This procedure can also tessellate all the geometric objects in a geometry table or view and add the tiles to the spatial index table.

Use this procedure to tessellate the geometries into fixed-size tiles.

### Syntax

```
SDO_ADMIN.POPULATE_INDEX_FIXED (layername, tile_size, [synch_flag,] [sdo_tile_flag,] [sdo_maxcode_flag])
```

### Keywords and Parameters

<i>layername</i>	Specifies the name of the data set layer. The <i>layername</i> is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
<i>tile_size</i>	Specifies the number of tessellations required to achieve the desired tile size (see the Usage Notes.) Data type is INTEGER.
<i>synch_flag</i>	Specifies whether to tessellate every geometric object in the geometry table, or only those that do not have corresponding entries in the spatial index table. If TRUE, only those geometric objects in the geometry table that do not have any corresponding tiles in the spatial index table are tessellated. If FALSE, all the geometric objects in the geometry table are tessellated and new tiles are simply added to the spatial index table. Default value is TRUE. Data type is BOOLEAN.
<i>sdo_tile_flag</i>	For internal use only. Not supported in this release. Default value is FALSE.

*sdo\_maxcode* Specifies whether or not the SDO\_MAXCODE column is populated. If  
*\_flag* TRUE, SDO\_MAXCODE is populated. If FALSE, the column is not popu-  
 lated.  
 Default value is TRUE.  
 Data type is BOOLEAN.

## Usage Notes

---



---

**Note:** This procedure is likely to be removed in a future release of Spatial Cartridge.

---



---

Consider the following when using this procedure:

- The SQL CREATE TABLE statement is used to create the spatial index table, <layername>\_SDOINDEX, prior to calling this procedure.
- The layer is tessellated into equal-sized tiles based on the number passed in the *tile\_size* parameter. The value of *tile\_size* specifies how many times to tessellate the layer. See Section 2.3.2, “Spatial Indexing with Fixed-Size Tiles”.
- For performance reasons, set the *synch\_flag* to FALSE when the spatial index table contains zero rows.
- For performance reasons, create an index on the SDO\_GID column in the <layername>\_SDOGEOM table before calling this routine.
- SDO\_ADMIN.POPULATE\_INDEX\_FIXED( ) behaves similarly to the CREATE INDEX statement in SQL. An implicit commit is executed after the procedure is called.
- SDO\_ADMIN.POPULATE\_INDEX\_FIXED( ) operates as a single transaction. To reduce the amount of rollback required to execute this procedure, you can write a routine that loops and calls SDO\_ADMIN.UPDATE\_INDEX\_FIXED( ) repeatedly. See Section A.1.1.1, “cr\_spatial\_index.sql Script” for more information.

Example 5–2 tessellates all the geometric objects in the LAYER1\_SDOGEOM table using up to 256 (4<sup>4</sup>) fixed-size tiles and adds the generated tiles to the LAYER1\_SDOINDEX table.

**Example 5-2**

```
SQL> EXECUTE SDO_ADMIN.POPULATE_INDEX_FIXED('layer1',4,FALSE,FALSE,FALSE);
```

**Related Topics**

- SDO\_ADMIN.UPDATE\_INDEX\_FIXED() procedure
- SDO\_TUNE.ESTIMATE\_TILING\_LEVEL() function



## SDO\_ADMIN.SDO\_CODE\_SIZE

### Purpose

This function determines the size that the SDO\_CODE column should be in the <layername>\_SDOINDEX table.

### Syntax

SDO\_ADMIN.SDO\_CODE\_SIZE (*layername*)

### Keywords and Parameters

<i>layername</i>	Specifies the name of the data set layer. Data type is VARCHAR2.
------------------	---

### Returns

This function returns the required size in bytes for the SDO\_CODE column.  
Data type is INTEGER.

### Usage Notes

The SDO\_CODE column is used to store the bit-interleaved cell ID of a tile that covers a geometry. The SDO\_MAXCODE column is SDO\_CODE padded out one place farther than the longest allowable code name for the index. Both columns are defined as RAW data types, with a maximum of 255 bytes. Use the SDO\_ADMIN.SDO\_CODE\_SIZE( ) function to fine-tune the size of the columns.

You should always set the SDO\_MAXCODE column to one greater than the SDO\_CODE column.

### Related Topics

None

## SDO\_ADMIN.UPDATE\_INDEX

### Purpose

This procedure tessellates a single geometric object in a geometry table or view and adds the tiles to the spatial index table. If the object already exists and has index entries, those entries are deleted and replaced by the newly generated tiles.

### Syntax

SDO\_ADMIN.UPDATE\_INDEX (*layername*, *GID*)

### Keywords and Parameters

<i>layername</i>	Specifies the name of the data set layer. The layername is used to construct the name of the geometry table. Data type is VARCHAR2.
<i>GID</i>	Specifies the geometric object identifier. Data type is NUMBER.

### Usage Notes

Consider the following when using this procedure:

- The <layername>\_SDOINDEX table must exist prior to calling this procedure. Use the SQL CREATE TABLE statement to create the spatial index table.
- For performance reasons, create an index on the SDO\_GID column in the <layername>\_SDOGEOM table before calling this routine.
- The values of the SDO\_LEVEL and SDO\_NUMTILES columns must be set in the <layername>\_SDOLAYER table before calling this procedure. This procedure generates either fixed or variable-sized tiles depending on values stored in the <layername>\_SDOLAYER table as follows:

SDO_LEVEL	SDO_NUMTILES	Action
NULL	NULL	Error
>= 1	NULL	Fixed-size tiling

SDO_LEVEL	SDO_NUMTILES	Action
>= 1	>= 1	Indexing with variable-sized tiles. The SDO_LEVEL column defines the partition bucket size. The SDO_NUMTILES column defines the number of tiles to generate per geometry. Note: variable-sized tiling is for experimentation purposes only.
NULL	>= 1	Not supported

- `SDO_ADMIN.UPDATE_INDEX()` does not perform an implicit commit after it executes and therefore the transaction can be rolled back.

Example 5-3 tessellates the polygon for geometry 25 and adds the generated tiles to the `LAYER1_SDOINDEX` table.

**Example 5-3**

```
SQL> EXECUTE SDO_ADMIN.UPDATE_INDEX('layer1', 25);
SQL> COMMIT;
```

## Related Topics

- `SDO_ADMIN.POPULATE_INDEX()` procedure

## SDO\_ADMIN.UPDATE\_INDEX\_FIXED

### Purpose

This procedure is provided for compatibility with Spatial Cartridge release 8.0.3 tables. This procedure has been replaced by enhanced features in the `SDO_ADMIN.UPDATE_INDEX()` procedure, and by supporting schema changes as shown in Section 1.4.

This procedure tessellates a single geometric object in a geometry table or view and adds the fixed-sized tiles to the spatial index table. By default, these tiles will replace existing ones for the same geometry; or optionally, existing tiles can be left alone.

### Syntax

```
SDO_ADMIN.UPDATE_INDEX_FIXED (layername, GID, tile_size, [replace_flag,] [sdo_tile_flag]  
[sdo_maxcode_flag])
```

### Keywords and Parameters

<i>layername</i>	Specifies the name of the data set layer. The <i>layername</i> is used to construct the name of the geometry table. Data type is VARCHAR2.
<i>GID</i>	Specifies the geometric object identifier. Data type is NUMBER.
<i>tile_size</i>	Specifies the number of tessellations required to achieve the desired fixed-size tiles. Each tessellation subdivides the tiles from the previous level into four smaller tiles. Data type is INTEGER.
<i>replace_flag</i>	Specifies whether or not to delete tiles for the <i>GID</i> before adding new ones. If TRUE, tiles are deleted prior to inserting new entries into the spatial index table. If FALSE, new tiles are simply added to the spatial index table. Default value is TRUE. Data type is BOOLEAN.
<i>sdo_tile_flag</i>	For internal use only. Not supported in this release. Default value is FALSE. Data type is BOOLEAN.

*sdo\_maxcode\_flag* Specifies whether or not the SDO\_MAXCODE column is populated. If TRUE, SDO\_MAXCODE is populated. If FALSE, the column is not populated.  
Default value is TRUE.  
Data type is BOOLEAN.

## Usage Notes

---

---

**Note:** This procedure is likely to be removed in a future release of Spatial Cartridge.

---

---

Consider the following when using this procedure:

- For performance reasons, set the *replace\_flag* to FALSE when the spatial index table contains no entries for the specified GID.
- For performance reasons, create an index on the SDO\_GID column in the <layername>\_SDOGEOM table before calling this procedure.
- SDO\_ADMIN.UPDATE\_INDEX\_FIXED( ) does not perform an implicit commit after it executes and therefore this transaction can be rolled back.

Example 5–4 tessellates the polygon for geometry 25 and adds the generated tiles to the LAYER1\_SDOINDEX table.

### **Example 5–4**

```
SQL> EXECUTE SDO_ADMIN.UPDATE_INDEX_FIXED ( 'layer1' , 25 , 4 , FALSE , FALSE , FALSE ) ;
```

## Related Topics

- SDO\_ADMIN.POPULATE\_INDEX\_FIXED() procedure
- SDO\_TUNE.ESTIMATE\_TILING\_LEVEL() function

## SDO\_ADMIN.VERIFY\_LAYER

---

### Purpose

This procedure checks for the existence of the geometry and spatial index tables.

### Syntax

```
SDO_ADMIN.VERIFY_LAYER (layername, [maxtiles])
```

### Keywords and Parameters

- layername* Specifies the name of the data set layer. The *layername* is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
- maxtiles* For internal use only. Not supported in this release.

### Usage Notes

If this procedure does not find the geometry and spatial index tables, it generates the following error: SDO 13113 (Oracle table does not exist)

Example 5-5 verifies the LAYER1 data set layer:

#### **Example 5-5**

```
SQL> EXECUTE SDO_ADMIN.VERIFY_LAYER('layer1');
```

### Related Topics

None

---

## Partitioned Point Data Procedures

Table 5–2 lists the procedures that can be used with partitioned point data. These procedures are neither required nor compatible with the geometry-based data format.

**Table 5–2 Administrative Procedures for Partitioned Point Data**

Procedure	Data Structure	Description
SDO_ADMIN.ALTER_HIGH_WATER_MARK	Partitioned points	Alters the high-water mark of a partitioned table.
SDO_ADMIN.DROP_PARTITION_INFO	Partitioned points	Removes a partitioned table.
SDO_ADMIN.PARTITION	Partitioned points	Places data into partition tables.
SDO_ADMIN.PROPAGATE_GRANTS	Partitioned points	Propagates the grants on the registered underlying table to the various partitions.
SDO_ADMIN.REGISTER_PARTITION_INFO	Partitioned points	Creates a partitioned spatial table.
SDO_ADMIN.REPARTITION	Partitioned points	Reorganizes a table based on the sorted values of the data contained within it.
SDO_ADMIN.VERIFY_PARTITIONS	Partitioned points	Checks for the existence of a table.

Also see Appendix A, “Sample SQL Scripts and Tuning Tips” for additional administrative tools useful for working with partitioned point data.

## SDO\_ADMIN.ALTER\_HIGH\_WATER\_MARK

### Purpose

This procedure alters the high-water mark of a partitioned spatial table. The high-water mark defines how many records can be stored in a partition before it subdivides. The table must exist and be registered in the Spatial Cartridge data dictionary.

This procedure is for use only with partitioned point data.

### Syntax

SDO\_ADMIN.ALTER\_HIGH\_WATER\_MARK (*tablename*, *high\_water\_mark*)

### Keywords and Parameters

<i>tablename</i>	Specifies the name of the partitioned table. Data type is VARCHAR2.
<i>high_water_mark</i>	Specifies the new high water mark for the table. Data type is INTEGER.

### Usage Notes

None

Example 5–6 changes the high-water mark to 5000 records for the TABLE1 partitioned spatial table.

#### **Example 5–6**

```
SQL> EXECUTE SDO_ADMIN.ALTER_HIGH_WATER_MARK('table1', 5000);
```

### Related Topics

- SDO\_ADMIN.REPARTITION() procedure
- altpart.sql sample SQL script file



## SDO\_ADMIN.DROP\_PARTITION\_INFO

### Purpose

This procedure removes a partitioned spatial table from the Spatial Cartridge data dictionary. The table must exist and must be registered in the Spatial Cartridge data dictionary.

This procedure is used only with partitioned point data.

### Syntax

SDO\_ADMIN.DROP\_PARTITION\_INFO (*tablename*)

### Keywords and Parameters

*tablename* Specifies the name of the partitioned table.  
Data type is VARCHAR2.

### Usage Notes

This procedure does not remove the spatial table and its associated partition tables from the user's schema. For a description of how to remove a partitioned spatial table from the user's schema, see the drppart.sql sample SQL script file described in Section A.1.2.2.

Example 5-7 removes the table1 table from the Spatial Cartridge data dictionary.

#### **Example 5-7**

```
SQL> EXECUTE SDO_ADMIN.DROP_PARTITION_INFO('table1');
```

### Related Topics

- drppart.sql sample SQL script file

## SDO\_ADMIN.PARTITION

### Purpose

This procedure places data into partition tables based on the sorted order of encoded dimensional values.

This procedure is used only with partitioned point data.

### Syntax

```
SDO_ADMIN.PARTITION (owner.source_table, tablename, parallel, guess , plummet_flag  
[,tablespace])
```

### Keywords and Parameters

<i>source_table</i>	Specifies the Oracle8 table or view of the table containing the partition key column. Data type is VARCHAR2.
<i>tablename</i>	Specifies the name of the table to partition. Data type is VARCHAR2.
<i>parallel</i>	Specifies the degree of parallelism for an operation on a single instance. Data type is INTEGER.
<i>guess</i>	Specifies the estimated largest common level of all the potential partitions to be created from data in the <i>source_table</i> . The common level of a partition is the number of levels of resolution of the common HHCODE for the partition. Data type is INTEGER.
<i>plummet_flag</i>	Specifies if the common HHCODE for all the potential partitions to be created from data in the <i>source_table</i> contains the maximum possible common level. If TRUE, the common HHCODE for each potential partition contains the maximum possible common level. If FALSE, the common HHCODE for each potential partition contains the minimum possible common level. Default value is FALSE. Data type is BOOLEAN.
<i>tablespace</i>	Specifies the tablespace in which the partitions should be created. Default is the tablespace of the underlying table.

## Usage Notes

Consider the following when using this procedure:

- The maximum size of the partition tables is determined by the high-water mark of the partitioned spatial table.
- To perform this procedure, first load the original data into an Oracle8 table using a utility such as SQL\*Loader. After the data is loaded, encode the data using the appropriate combination of Spatial Cartridge data conversion functions (see Chapter 9.) The encoded data is used as the partition key column. The partition key column is provided as either a column in the Oracle8 table or as a view of that table.
- For more information on specifying the degree of parallelism, see the *Oracle8 Server Tuning* manual.

Example 5–8 partitions the TABLE1 partitioned spatial table with data contained in the source1 table.

### **Example 5–8**

```
SQL> EXECUTE SDO_ADMIN.PARTITION('source1', 'table1', 1, 10, FALSE);
```

## Related Topics

- SDO\_ADMIN.REGISTER\_PARTITION\_INFO() procedure

## SDO\_ADMIN.PROPAGATE\_GRANTS

### Purpose

This procedure is used to propagate the grants on the underlying table to the partitions.

This procedure is used only with partitioned point data.

### Syntax

```
SDO_ADMIN.PROPAGATE_GRANTS (tablename)
```

### Keywords and Parameters

<i>tablename</i>	Specifies the name of the partitioned table. Data type is VARCHAR2.
------------------	--

### Usage Notes

This procedure is used after calls to `SDO_ADMIN.PARTITION()` or `SDO_ADMIN.REPARTITION()`. It must be called by the owner of the partition.

This procedure must be compiled prior to use. See Section A.1.2.3, “`sdogrant.sql` Script”.

Example 5–9 propagates grants from the TABLE1 partitioned spatial table.

#### **Example 5–9**

```
SQL> EXECUTE SDO_ADMIN.PROPAGATE_GRANTS('TABLE1');
```

### Related Topics

- `SDO_ADMIN.PARTITION()` procedure
- `SDO_ADMIN.REPARTITION()` procedure

## SDO\_ADMIN.REGISTER\_PARTITION\_INFO

### Purpose

This procedure creates a partitioned spatial table entry in the Spatial Cartridge data dictionary, and defines the partition key column and the high-water mark for the table.

This procedure is used only with partitioned point data.

### Syntax

SDO\_ADMIN.REGISTER\_PARTITION\_INFO (*tablename*, *column*, *high\_water\_mark*)

### Keywords and Parameters

<i>tablename</i>	Specifies the name of the partitioned table. Data type is VARCHAR2.
<i>column</i>	Specifies the name of the partition key column for the table. Data type is VARCHAR2.
<i>high_water_mark</i>	Specifies the number of records to store in a partition before the partition subdivides. Data type is INTEGER.

### Usage Notes

The SQL CREATE TABLE statement is used to create the partitioned spatial table, with the partition key column defined as RAW(255), prior to calling this procedure.

Example 5-10 registers the TABLE1 partitioned spatial table.

#### **Example 5-10**

```
SQL> EXECUTE SDO_ADMIN.REGISTER_PARTITION_INFO('table1',  
2> 'hhcolumn', 1000);
```

### Related Topics

- SDO\_ADMIN.PARTITION() procedure

## SDO\_ADMIN.REPARTITION

### Purpose

This procedure reorganizes a partitioned spatial table based on the sorted order of encoded dimensional values already contained in it. The table must exist and must be registered in the Spatial Cartridge data dictionary.

This procedure is used only with partitioned point data.

### Syntax

```
SDO_ADMIN.REPARTITION (tablename, parallel, [tablespace])
```

### Keywords and Parameters

<i>tablename</i>	Specifies the name of the partitioned table. Data type is VARCHAR2.
<i>parallel</i>	Specifies the degree of parallelism for an operation on a single instance. Data type is INTEGER.
<i>tablespace</i>	Specifies the name of the tablespace in which to create the partition. Data type is VARCHAR2.

### Usage Notes

Consider the following when using this procedure:

- The tablespace variable is optional. If you do not supply a tablespace name, the partitions are created in the same tablespace as the registered partition table.
- The maximum size of the reorganized partition tables is determined by the high-water mark of the partitioned spatial table.
- For more information on specifying the degree of parallelism, see the section on "Parallel Query Option," in the Oracle8 Server documentation.

Example 5–11 repartitions the table1 partitioned spatial table.

#### **Example 5–11**

```
SQL> EXECUTE SDO_ADMIN.REPARTITION('table1', 1);
```

**Related Topics**

- SDO\_ADMIN.ALTER\_HIGH\_WATER\_MARK() procedure

## SDO\_ADMIN.VERIFY\_PARTITIONS

### Purpose

This procedure checks if the partitioned spatial table exists, if it is registered in the Spatial Cartridge data dictionary, and if the partition key column exists as defined in the Spatial Cartridge data dictionary.

This procedure is used only with partitioned point data.

### Syntax

SDO\_ADMIN.VERIFY\_PARTITIONS (*tablename*)

### Keywords and Parameters

*tablename* Specifies the name of the table.  
Data type is VARCHAR2.

### Usage Notes

This procedure can generate the following errors depending on the results of the verification:

- SDO 13113 (Oracle table does not exist)
- SDO 13108 (spatial table not found)
- SDO 13111 (spatial table has no partition key defined)
- SDO 13129 (HHCODE column not found)

Example 5–12 verifies the TABLE1 partitioned spatial table:

**Example 5–12**

```
SQL> EXECUTE SDO_ADMIN.VERIFY_PARTITIONS('table1');
```

### Related Topics

- SDO\_ADMIN.REGISTER\_PARTITION\_INFO() procedure



---

# Tuning Functions

This chapter contains descriptions of the tuning functions and procedures shown in Table 6-1

**Table 6-1** *Tuning Functions and Procedures:*

Function/Procedure	Description
SDO_TUNE.ESTIMATE_TILING_LEVEL	Determines an appropriate tiling level for creating fixed-size index tiles.
SDO_TUNE.EXTENT_OF	Determines the minimum-bounding rectangle of the data in a layer.

---

## SDO\_TUNE.ESTIMATE\_TILING\_LEVEL

### Purpose

This function estimates the appropriate tiling level to use when indexing with fixed-size tiles.

### Syntax

SDO\_TUNE.ESTIMATE\_TILING\_LEVEL (*layername*, *maxtiles*, *type\_of\_estimate*)

### Keywords and Parameters

<i>layername</i>	Specifies the name of the dataset layer to examine. Data type is VARCHAR2.
<i>maxtiles</i>	Specifies the maximum number of tiles that can be used to index the rectangle defined by the <i>type_of_estimate</i> parameter. Data type is INTEGER.
<i>type_of_estimate</i>	Indicates by keyword one of three different models. Specify the type of estimate with one of the following keywords: <ul style="list-style-type: none"><li>• LAYER_EXTENT -- Use the rectangle defined by your coordinate system.</li><li>• ALL_GID_EXTENT -- Use the minimum-bounding rectangle that encompasses all the geometric objects within the layer.</li><li>• AVG_GID_EXTENT -- Use a rectangle representing the average size of the individual geometries within the layer. This option performs the most analysis of the three types.</li></ul>

### Returns

The function returns an integer representing the level to use when creating a spatial index for the specified layer.

### Usage Notes

The SDO\_ADMIN.POPULATE\_INDEX() and SDO\_ADMIN.UPDATE\_INDEX() procedures are used to create or update the spatial index using fixed-size or variable-

sized tiles. Store the value returned by the `SDO_TUNE.ESTIMATE_TILING_LEVEL()` function in the `SDO_LEVEL` column of the `<layername>_SDOLAYER` table prior to building the spatial index.

The `maxtiles` parameter specifies the maximum number of tiles that should be used to define a grid covering the rectangular extent of interest. This extent could be:

- the extent defined in the `<layername>_SDODIM` table which defines the bounds of the coordinate system
- the extent defined by the minimum and maximum coordinates for the given dataset (as returned by the `SDO_TUNE.EXTENT_OF()` procedure)
- the extent defined by computing the average bounds of the objects in the `<layername>_SDOGEOM` table

The code shown in Example 6-1 generates a recommendation based on the extent of the defined coordinate system (-90 to +90 latitude and -180 to +180 longitude). This example returns a level whose tiles are not smaller than one degree cells.

**Example 6-1 Recommended Tile Level for One-Degree Lat/Lon Cells**

```
set serveroutput on
declare
    tiling_level integer;
begin
    tiling_level := mdsys.sdo_tune.estimate_tiling_level('WORLD_CITIES',
        360*180, 'LAYER_EXTENT');
    dbms_output.put_line('VALUE is '|| tiling_level);
end;
```

For many applications, however, it is more effective to call the `SDO_TUNE.ESTIMATE_TILING_LEVEL()` function using the `ALL_GID_EXTENT` estimate type. In Example 6-2, assume the dataset consists of block groups for San Francisco and that the `<layername>_SDODIM` table defines the extent to be one that covers all of California. Because the dataset is localized to a small subregion of this extent, `ALL_GID_EXTENT` is the appropriate estimate type. The recommended tiling level in this case will be such that at most 10,000 tiles will be required to completely cover the extent of San Francisco block groups.

**Example 6-2 Recommended Tile Level Based on the GIDs of All Geometries**

```
set serveroutput on
declare
    tiling_level integer;
```

```
begin
  tiling_level:= mdsys.sdo_tune.estimate_tiling_level('SF_BLOCK_GROUPS',
    10000, 'ALL_GID_EXTENT');
  dbms_output.put_line('VALUE is' ,|| tiling_level);
end;
```

The third type of estimate helps determine the tiling level that should be used such that on average, the maxtiles parameter defines the number of tiles to cover the extent of a single geometry in the layer. This estimate type requires the most computation of the three because the bounding box of every geometry is used in calculating the average extent. In Example 6–3, eight tiles on average are used to cover any block group in San Francisco.

**Example 6–3 Recommended Tile Level Based on Average Extent of All Geometries**

```
set serveroutput on
declare
  tiling_level integer;
begin
  tiling_level := mdsys.sdo_tune.estimate_tiling_level('SF_BLOCK_GROUPS', 8,
    'AVG_GID_EXTENT');
  dbms_output.put_line('Tiling level value is ' || tiling_level);
end;
```

## Related Topics

- SDO\_ADMIN.POPULATE\_INDEX
- SDO\_ADMIN.UPDATE\_INDEX
- SDO\_TUNE.EXTENT\_OF
- Section A.2.2, “Understanding the Tiling Level”
- Section A.2.8, “Visualizing the Spatial Index (Drawing Tiles)”

---

## SDO\_TUNE.EXTENT\_OF

### Purpose

This procedure determines the extent of all geometries in a layer.

### Syntax

```
SDO_TUNE.EXTENT_OF (layername, min_X, max_X, min_Y, max_Y)
```

### Keywords and Parameters

<i>layername</i>	Specifies the name of the dataset layer. The layer-name is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
<i>min_X</i>	Minimum X value of the bounding rectangle. Data type is number.
<i>max_X</i>	Maximum X value of the bounding rectangle. Data type is number.
<i>min_Y</i>	Minimum Y value of the bounding rectangle. Data type is number.
<i>max_Y</i>	Maximum Y value of the bounding rectangle. Data type is number.

### Returns

This procedure returns the coordinates of the minimum-bounding rectangle for all geometric data in a layer. The data type is NUMBER for the four return values.

### Usage Notes

None

### Related Topics

- SDO\_TUNE.ESTIMATE\_TILING\_LEVEL() function



---

# Geometry Functions

This chapter contains descriptions of the geometric functions and procedures shown in Table 7-1.:

**Table 7-1 Geometric Functions and Procedures**

<b>Function/Procedure</b>	<b>Description</b>
SDO_GEOM.ADD_NODES	Stores points in the geometry table.
SDO_GEOM.INIT_ELEMENT	Initializes space in the geometry table for a new object.
SDO_GEOM.INTERACT	Determines if two objects are disjoint.
SDO_GEOM.RELATE	Determines how two objects interact.
SDO_GEOM.VALIDATE_GEOMETRY	Determines if a geometry is valid.

## SDO\_GEOM.ADD\_NODES

### Purpose

This procedure stores coordinate geometry points into the SDOGEOM table.

### Syntax

```
SDO_GEOM.ADD_NODES (layername, SDO_GID, SDO_ESEQ, SDO_ETYPE, X-ord1, Y-ord1[, ..., X125, Y125])
```

### Keywords and Parameters

<i>layername</i>	Specifies the name of the data set layer. The layer-name is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
<i>SDO_GID</i>	Specifies the unique geometric object identifier. Data type is NUMBER.
<i>SDO_ESEQ</i>	Specifies the element sequence number. The <i>eseq</i> is unique within an <i>SDO_GID</i> . Data type is NUMBER.
<i>SDO_ETYPE</i>	Specifies the type of geometric element. Data type is INTEGER, corresponding to the following constants:  <ol style="list-style-type: none"><li>1 SDO_GEOM.POINT_TYPE</li><li>2 SDO_GEOM.LINESTRING_TYPE</li><li>3 SDO_GEOM.POLYGON_TYPE</li></ol>
<i>X ordinateN</i> , <i>Y ordinateN</i>	Specifies the X and Y values of a vertex (coordinate pair) in a geometry. Up to 125 pairs may be added in a single call. Data type is NUMBER.

### Usage Notes

Consider the following when using this procedure:

- Use the SQL CREATE TABLE statement to create the geometry table, <layername>\_SDOGEOM, before calling this procedure.



- Prior to calling this procedure, call the `SDO_GEOM.INIT_ELEMENT()` function to initialize the geometry element and retrieve the element sequence number (`SDO_ESEQ`).
- Close a polygon by providing the coordinates of the first vertex as the last vertex.
- Call this procedure iteratively with the same `GID` to add coordinates to a geometric object. You can add up to 125 coordinate pairs on each call, and there is no limit to how many times you can add more vertices.
- This procedure cannot be used when `<layername>_SDOGEOM` is a view. For layer objects created as views, you need to explicitly insert new geometries, as opposed to using `SDO_GEOM.INIT_ELEMENT()` and `SDO_GEOM.ADD_NODES()`.

Example 7-1 adds a polygon element for geometry 25 in the `LAYER1` data set. The polygon is a square.

**Example 7-1**

```
SQL> EXECUTE SDO_GEOM.ADD_NODES ('LAYER1', 25,  
2> SDO_GEOM.POLYGON_TYPE,  
3> 3,3, 7,3,  
4> 7,7, 3,7,  
5> 3,3);
```

## Related Topics

- `SDO_GEOM.INIT_ELEMENT()` function

## SDO\_GEOM.INIT\_ELEMENT

### Purpose

This function initializes elements in the SDOINFO table or view for a new geometry element.

### Syntax

```
SDO_GEOM.INIT_ELEMENT (layername, SDO_GID)
```

### Keywords and Parameters

<i>layername</i>	Specifies the name of the data set layer. The layer-name is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
<i>SDO_GID</i>	Specifies the geometric object identifier. Data type is NUMBER.

### Returns

This function returns the next element sequence number. The data type is INTEGER.

### Usage Notes

This function initializes the element to be stored, but does not actually insert coordinates into the SDOGEOM table. The `SDO_GEOM.ADD_NODES ( )` procedure is used to insert associated coordinate data.

For layer objects created as views, you need to explicitly insert new geometries, as opposed to using `SDO_GEOM.INIT_ELEMENT ( )` and `SDO_GEOM.ADD_NODES ( )`.

### Related Topics

- `SDO_GEOM.ADD_NODES()` procedure

---

## SDO\_GEOM.INTERACT

### Purpose

This function determines if two geometry objects interact.

### Syntax

SDO\_GEOM.INTERACT (*layername1*, *SDO\_GID1*, [*layername2*,] *SDO\_GID2*)

SDO\_GEOM.INTERACT (*layername1*, *SDO\_GID1*, *X\_tolerance*, *Y\_tolerance*, *SDO\_ETYPE*,  
*num\_ordinates*, *X\_ordinate1*, *Y\_ordinate1* [, ..., *Xn*, *Yn*]  
[,*SDO\_ETYPE*, *num\_ordinates*, *X\_ordinate1*, *Y\_ordinate1* [,...*Xn*,*Yn*]])

### Keywords and Parameters

<i>layername1</i> , <i>layername2</i>	Specifies the name of the data set layer. The layername is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
<i>SDO_GID1</i> , <i>SDO_GID2</i>	Specifies the geometric object identifier. Data type is NUMBER.
<i>X_tolerance</i> , <i>Y_tolerance</i>	Specifies the distance two points can be apart and still be considered the same due to rounding errors. Tolerance must be greater than zero. If you want zero tolerance, enter a number such as 0.000005, where the number of zeroes to the right of the decimal point matches the precision of your data. Data type is NUMBER.
<i>SDO_ETYPE</i>	Specifies the type of geometry object. Data type is INTEGER, corresponding to the following constants: <ul style="list-style-type: none"> <li>1 SDO_GEOM.POINT_TYPE</li> <li>2 SDO_GEOM.LINestring_TYPE</li> <li>3 SDO_GEOM.POLYGON_TYPE</li> </ul>

<i>num_ordinates</i>	Specifies the number of ordinates for this element. Data type is NUMBER.
<i>X_ordinateN</i> , <i>Y_ordinateN</i>	Specifies the X and Y values of a vertex (coordinate pair) in a geometry. Data type is NUMBER.

## Returns

This function returns TRUE if the first and second objects interact with each other and are not disjoint. The data type is VARCHAR2.

## Usage Notes

Use the first form of the function to test two stored geometry objects.

Use the second form of the function to compare a stored object against a user-defined object. You can specify up to 123 vertices for a single element geometry. If the geometry has multiple elements, the total number of arguments passed, including SDO\_ETYPE, num\_ordinates, and the list of vertex coordinates cannot exceed 250 values.

## Related Topics

- SDO\_GEOM.RELATE() function

---

## SDO\_GEOM.RELATE

### Purpose

This function examines two geometry objects to determine their spatial relationship.

### Syntax

SDO\_GEOM.RELATE (*layername1*, *SDO\_GID1*, *mask*, [*layername2*,] *SDO\_GID2*)

SDO\_GEOM.RELATE (*layername1*, *SDO\_GID1*, *mask*, *X\_tolerance*, *Y\_tolerance*, *SDO\_ETYPE*, *num\_ordinates*, *X\_ordinate1*, *Y\_ordinate1* [, ..., *Xn*, *Yn*] [, *SDO\_ETYPE*, *num\_ordinates*, *X\_ordinate1*, *Y\_ordinate1* [, ..., *Xn*, *Yn*]])

### Keywords and Parameters

<i>layername1</i> , <i>layername2</i>	Specifies the name of the data set layer. The layer-name is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
<i>SDO_GID1</i> , <i>SDO_GID2</i>	Specifies the geometry object identifier. Data type is NUMBER.
<i>mask</i>	Specifies the relationship mask. See the list of keywords in the Usage Notes.
<i>X_tolerance</i> , <i>Y_tolerance</i>	Specifies the distance two points can be apart and still be considered the same due to rounding errors. Tolerance must be greater than zero. If you want zero tolerance, enter a number such as 0.000005, where the number of zeroes to the right of the decimal point matches the precision of your data. Data type is NUMBER.
<i>SDO_ETYPE</i>	Specifies the type of geometry element. Data type is INTEGER, corresponding to the following constants: <ol style="list-style-type: none"> <li>1 SDO_GEOM.POINT_TYPE</li> <li>2 SDO_GEOM.LINestring_TYPE</li> <li>3 SDO_GEOM.POLYGON_TYPE</li> </ol>

<i>num_ordinates</i>	Specifies the number of ordinates for this element. Data type is NUMBER.
<i>X_ordinateN</i> , <i>Y_ordinateN</i>	Specifies the X and Y values of a vertex (coordinate pair) in a geometry. Data type is NUMBER.

## Returns

The `SDO_GEOM.RELATE()` function can return three types of answers:

1. If you pass a mask listing one or more relationships, the function returns the names of the relationships if all of them are true. If one or more relationships are false, the procedure returns `FALSE`.
2. If you pass the `DETERMINE` keyword in the mask, the function returns the one relationship keyword that best matches the geometries.
3. If you pass the `ANYINTERACT` keyword in the mask, the function returns `TRUE` if the two geometries are not disjoint. This is equivalent to the `SDO_GEOM.INTERACT()` function.

The data type is `VARCHAR2`.

## Usage Notes

Use the first form of the function to examine two stored geometry objects.

Use the second form of the function to compare a stored object against a user-defined object. You can specify up to 123 vertices for a single-element geometry. If the geometry has multiple elements, the total number of arguments passed, including `SDO_ETYPE`, `num_ordinates`, and the list of vertex coordinates cannot exceed 250 values.

The following relationships can be tested:

- `ANYINTERACT` - Returns `TRUE` if the objects are not disjoint.
- `CONTAINS` - Returns `TRUE` if the second object is entirely within the first object and the object boundaries do not touch.
- `COVEREDBY` - Returns `TRUE` if the first object is entirely within the second object and the object boundaries touch at one or more points.
- `COVERS` - Returns `TRUE` if the second object is entirely within the first object and the boundaries touch in one or more places.

- **DISJOINT** - Returns TRUE if the objects have no common boundary or interior points.
- **EQUAL** - Returns TRUE if the objects share every point of their boundaries and interior, including any holes in the objects.
- **INSIDE** - Returns TRUE if the first object is entirely within the second object and the object boundaries do not touch.
- **OVERLAPBDYDISJOINT** - Returns TRUE if the objects overlap, but their boundaries do not interact.
- **OVERLAPBDYINTERSECT** - Returns TRUE if the object overlap, and their boundaries intersect in one or more places.
- **TOUCH** - Returns TRUE if the two objects share a common boundary point, but no interior points.

Mask values can be combined using a logical OR. For example, 'INSIDE + TOUCH' returns TRUE if the objects pass either test.

## Related Topics

- `SDO_GEOM.INTERACT()` function

## SDO\_GEOM.VALIDATE\_GEOMETRY

### Purpose

This function provides a consistency check for valid geometry types. The function checks the representation of the geometry from the tables against the element definitions.

### Syntax

SDO\_GEOM.VALIDATE\_GEOMETRY (*layername*,*SDO\_GID*)

### Keywords and Parameters

<i>layername</i>	Specifies the name of the data set layer. The <i>layername</i> is used to construct the name of the geometry and spatial index tables. Data type is VARCHAR2.
<i>SDO_GID</i>	Specifies the geometric object identifier. Data type is NUMBER.

### Returns

This function returns TRUE if the geometry is valid. The data type is VARCHAR2.

### Usage Notes

This function checks for the following:

- polygons have at least three points and must be closed
- line strings must have at least two points
- when and SDO\_ESEQ spans multiple rows, the last point of the previous row is the first point on the next row

### Related Topics

None



---

## Window Functions

If a query window does not already exist in the database, you must first insert it and create and index for it. The SDO\_WINDOW functions are used to create temporary geometry objects to be used in comparisons with stored geometries. You can create query windows with any number of coordinates.

Because not all Oracle users may have insert privileges, the SDO\_WINDOW package is not automatically installed when you install Spatial Cartridge. This allows a DBA to control the schema under which this package operates. Choose an Oracle user who has insert privilege and compile the SDO\_WINDOW package under that user. For example, you could choose the mdsys Oracle user:

```
% sqlplus mdsys/password
SQL> @$ORACLE_HOME/md/admin/sdowin.sql
SQL> @$ORACLE_HOME/md/admin/prvtwin.plb
```

This chapter contains descriptions of the window functions listed in Table 8–1.

**Table 8–1 Window Functions**

Function	Description
SDO_WINDOW.BUILD_WINDOW	Builds a query window geometry object.
SDO_WINDOW.BUILD_WINDOW_FIXED	Builds a query window using fixed-size tiles.
SDO_WINDOW.CLEAN_WINDOW	Drops the tables used for a query window.
SDO_WINDOW.CREATE_WINDOW_LAYER	Creates the tables needed for a query window layer.

## SDO\_WINDOW.BUILD\_WINDOW

---

### Purpose

This function builds the window for the query and returns an SDO\_GID that serves as a handle. The window is tessellated into variable-sized tiles.

### Syntax

```
SDO_WINDOW.BUILD_WINDOW(comp_name, layername, SDO_ETYPE, SDO_NUMTILES, X1, Y1,  
[...Xn, Yn])
```

### Keywords and Parameters

<i>comp_name</i>	Specifies the name of the user who compiled this package. This user must have appropriate privileges to read and write into the database. Data type is VARCHAR2.
<i>layername</i>	Specifies the name of the window layer into which the coordinates will be inserted. Data type is VARCHAR2.
<i>SDO_ETYPE</i>	Specifies the type of geometry element. Data type is INTEGER, corresponding to the following constants: <ul style="list-style-type: none"><li>1 SDO_GEOM.POINT_TYPE</li><li>2 SDO_GEOM.LINESTRING_TYPE</li><li>3 SDO_GEOM.POLYGON_TYPE</li></ul>
<i>SDO_NUMTILES</i>	Data type is NUMBER, value must be NULL.
<i>X ordinateN</i> , <i>Y ordinateN</i>	Specifies the X and Y values of a vertex (coordinate pair) in a geometry. Up to 125 pairs may be added in a single call. Data type is NUMBER.

### Returns

This function returns the SDO\_GID of the new geometry. The data type is NUMBER.

## Usage Notes

This function inserts the coordinates into the <layername>\_SDOGEOM table, tessellates the geometry (creates the index), and returns a unique SDO\_GID corresponding to the geometry.

You do not need special privileges to execute this function. However, the user who compiles it does need appropriate privileges to read and write into the database.

When working with Spatial Cartridge release 8.0.3 tables, the SDO\_NUMTILES parameter indicates the number of tiles into which the window should be tessellated. For release 8.0.4, the function reads that information from the <layername>\_SDOLAYER table.

## Related Topics

- SDO\_WINDOW.BUILD\_WINDOW\_FIXED() function

## SDO\_WINDOW.BUILD\_WINDOW\_FIXED

### Purpose

This function builds the window for the query and returns an SDO\_GID that serves as a handle. The window is tessellated into fixed-size tiles.

### Syntax

SDO\_WINDOW.BUILD\_WINDOW\_FIXED (*comp\_name*, *layername*, *SDO\_ETYPE*, *SDO\_TILESIZE*, *X1*, *Y1*, [...*Xn*, *Yn*])

### Keywords and Parameters

<i>comp_name</i>	Specifies the name of the user who compiled this package. This user must have appropriate privileges to read and write into the database. Data type is VARCHAR2.
<i>layername</i>	Specifies the name of the window layer into which the coordinates will be inserted. Data type is VARCHAR2.
<i>SDO_ETYPE</i>	Specifies the type of geometry element. Data type is INTEGER, corresponding to the following constants: <ul style="list-style-type: none"><li>1 SDO_GEOM.POINT_TYPE</li><li>2 SDO_GEOM.LINESTRING_TYPE</li><li>3 SDO_GEOM.POLYGON_TYPE</li></ul>
<i>SDO_TILESIZE</i>	Specifies the number of tessellations required to achieve the desired fixed-size tiles. Data type is NUMBER.
<i>X ordinateN</i> , <i>Y ordinateN</i>	Specifies the X and Y values of a vertex (coordinate pair) in a geometry. Up to 125 pairs may be added in a single call. Data type is NUMBER.

### Returns

This function returns the SDO\_GID of the new geometry. Data type is NUMBER.

## Usage Notes

---

---

**Note:** This procedure is likely to be replaced in a future release of Spatial Cartridge.

---

---

This function inserts the coordinates into the <layername>\_SDOGEOM table, tessellates the geometry (creates the index), and returns a unique SDO\_GID corresponding to the geometry.

You do not need special privileges to execute this function. However, the user who compiles it does need appropriate privileges to read and write into the database.

Query SDO\_LEVEL from the <layername>\_SDOLAYER table to pass the correct SDO\_TILE\_SIZE value to this function.

## Related Topics

None

---

## SDO\_WINDOW.CLEAN\_WINDOW

### Purpose

This procedure drops the four tables created in the layer for the query window.

### Syntax

```
SDO_WINDOW.CLEAN_WINDOW (layername);
```

### Keywords and Parameters

<i>layername</i>	Specifies the name of the window layer that needs to be dropped. Data type is VARCHAR2.
------------------	--

### Usage Notes

Typically, you would build a layer once, and then build multiple windows and perform multiple queries using that layer. After finishing all queries, you can execute the `SDO_WINDOW.CLEAN_WINDOW( )` procedure to drop the tables.

### Related Topics

None

---

## SDO\_WINDOW.CREATE\_WINDOW\_LAYER

### Purpose

This procedure creates the necessary tables that constitute a layer used for defining a query window.

### Syntax

```
SDO_WINDOW.CREATE_WINDOW_LAYER (layername, SDO_LEVEL, SDO_NUMTILES,  
SDO_DIMNUM1, SDO_LB1, SDO_UB1, SDO_TOLERANCE1, SDO_DIMNAME1, SDO_DIMNUM2,  
SDO_LB2, SDO_UB2, SDO_TOLERANCE2, SDO_DIMNAME2)
```

### Keywords and Parameters

<i>layername</i>	Specifies the name of the window layer to be created. The <i>layername</i> is used to construct the four tables associated with the layer. Data type is VARCHAR2.
SDO_LEVEL	Specifies the number of times the layer should be tessellated during the indexing phase. Data type is INTEGER.
SDO_NUMTILES	Specifies the number of tiles to generate during indexing. Data type is INTEGER.
SDO_DIMNUM1, SDO_DIMNUM2	Specifies the number of the dimension, starting with 1. Data type is NUMBER.
SDO_LB1, SDO_UB1, SDO_LB2, SDO_UB2	Specifies the lower and upper bounds of this dimension. Data type is NUMBER.
SDO_TOLERANCE1, SDO_TOLERANCE2	Specifies the allowable variance of ordinate values within each dimension. Data type is NUMBER.
SDO_DIMNAME1, SDO_DIMNAME2	Specifies the name of the dimension. Data type is VARCHAR2.

## Usage Notes

Because the <layername>\_SDODIM table is initialized with the dimension and the bound information, only those queries that are in the same dimension should be queried against this layer. If you wish to issue a query with respect to a different dimension, you must create a new layer.

## Related Topics

None



---

## Partitioned Point Data Functions

Spatial Cartridge has undergone an architectural change, beginning with the 7.3.3 release. The emphasis on partitioned tables has been replaced by the improved spatial indexing features.

The functions described in this chapter are not required for creating or maintaining a spatial database, however, they are provided for convenience in working with legacy data in partitioned point data tables. They are used with SQL SELECT, INSERT, UPDATE, and DELETE statements to perform the following:

- generate dimensions from bounded, hierarchical, or date data values
- encode and decode dimensions
- retrieve bounded, hierarchical, or date data values from dimensions

When using these functions in basic SQL statements, use the form: SDO\_<function>. When using the functions inside a PL/SQL block, use a period (.) instead of the underscore.

This chapter contains descriptions of the spatial functions listed in Table 9-1.

**Table 9-1** *Partitioned Point Data Functions*

Function	Purpose
SDO_BVALUETODIM	Creates a dimension from bounded data values.
SDO_COMPARE	Evaluates the relationship between two objects described by HHCODEs.
SDO_DATETODIM	Creates a dimension from an Oracle DATE data type.
SDO_DECODE	Extracts a single dimension from an HHCODE.

---

**Table 9–1 Partitioned Point Data Functions**

<b>Function</b>	<b>Purpose</b>
SDO_ENCODE	Creates an HHCODE by combining dimensions to describe an area or point.
SDO_TO_BVALUE	Extracts a bounded data value from a dimension.
SDO_TO_DATE	Extracts an Oracle DATE data type from a dimension.

Additional functions that support partitioned point data can be found in Chapter 5, “Administrative Procedures” and Appendix A, “Sample SQL Scripts and Tuning Tips”.

---

## SDO\_BVALUETODIM

### Purpose

This function creates a dimension from a bounded value, which is a value contained in a set of values expressed as a lower boundary and an upper boundary.

### Syntax

SDO\_BVALUETODIM (*value*, *lower\_boundary*, *upper\_boundary*, *decimal\_scale*)

### Keywords and Parameters

<i>value</i>	Specifies the value for the particular dimension. Data type is NUMBER.
<i>lower_boundary</i>	Specifies the lower boundary of the dimension range. Data type is NUMBER.
<i>upper_boundary</i>	Specifies the upper boundary of the dimension range. Data type is NUMBER.
<i>decimal_scale</i>	Specifies the number of digits to the right of the decimal point. Data type is NUMBER.

### Returns

This function returns a dimension. The data type is RAW.

### Usage Notes

Example 9–1 shows the SDO\_BVALUETODIM() function.

#### **Example 9–1**

```
SQL> INSERT INTO sourcetable1(SAMPLENAME,DATA_PT)
2> VALUES ('SAMPLE1',SDO_ENCODE(SDO_BVALUETODIM(10,-100,100,7),
3> SDO_BVALUETODIM(20,-100,100,7)));
```

### Related Topics

- SDO\_ENCODE() function
- SDO\_TO\_BVALUE() function

## SDO\_COMPARE

---

### Purpose

This function evaluates the relationship between an area or point described by an HHCODE and another HHCODE, or a range of HHCODEs expressed as an upper bound and lower bound.

### Syntax

```
SDO_COMPARE (hhcode_expression, {hhcode_expression |  
lower_bound_HHCODE,upper_bound_HHCODE})
```

### Keywords and Parameters

<i>hhcode_expression</i>	Specifies an expression that evaluates to an HHCODE. Data type is RAW.
<i>lower_bound_HHCODE</i>	Specifies the lower bound HHCODE expression. Data type is RAW.
<i>upper_bound_HHCODE</i>	Specifies the upper bound HHCODE expression. Data type is RAW.

### Returns

This function returns one of the following keywords:

- ENCLOSES
- EQUAL
- INSIDE
- OUTSIDE
- OVERLAP

The data type is VARCHAR2.

### Usage Notes

Example 9-2 selects all points that fall within the given multidimensional range.

**Example 9–2**

```
SQL> SELECT SDO_GID FROM layer1_SDOINDEX WHERE  
2> SDO_COMPARE(SDO_MAXCODE,  
3> SDO_ENCODE(5,5),  
4> SDO_ENCODE(25,25))='INSIDE';
```

Example 9–3 selects GIDs based on interaction between their spatial index tiles.

**Example 9–3**

```
SQL> SELECT SDO_GID FROM layer1_SDOINDEX A, layer2_SDOINDEX B  
2> WHERE SDO_COMPARE(A.SDO_CODE,B.SDO_CODE) != 'OUTSIDE';
```

**Related Topics**

- SDO\_GEOM.INTERACT() function
- SDO\_GEOM.RELATE() function

## SDO\_DATETODIM

### Purpose

This function creates a dimension from an Oracle DATE data type. The component number determines the level of resolution of the date in the dimension.

### Syntax

SDO\_DATETODIM (*date* [, *component*])

### Keywords and Parameters

<i>date</i>	Specifies the calendar date. Data type is DATE.
<i>component</i>	Specifies the level of resolution. The component number values are defined as follows:  <ol style="list-style-type: none"><li>1 accurate to year</li><li>2 accurate to month</li><li>3 accurate to day</li><li>4 accurate to hour</li><li>5 accurate to minute</li><li>6 accurate to second</li></ol> The default value is 6. Data type is INTEGER.

### Returns

This function returns a dimension. The data type is RAW.

### Usage Notes

You must use a valid Oracle date format string.

Example 9–4 shows the SDO\_DATETODIM( ) function.

#### **Example 9–4**

```
SQL> INSERT INTO sourcetable1(SAMPLENAME,DATA_PT)
2> VALUES('SAMPLE1',SDO_ENCODE(SDO_DATETODIM(TO_DATE('19-Jul-96'),
3> SDO_EVALUETODIM(100,-1000,1000,7)));
```

**Related Topics**

- SDO\_ENCODE() function
- SDO\_TO\_DATE() function

## SDO\_DECODE

### Purpose

This function extracts a single dimension from an HHCODE.

### Syntax

`SDO_DECODE (hhcode_expression, dimension_number)`

### Keywords and Parameters

<i>hhcode_expression</i>	Specifies an expression that evaluates to an HHCODE. Data type is RAW.
<i>dimension_number</i>	Specifies the dimension number to extract. Data type is INTEGER.

### Returns

This function returns a dimension. The data type is RAW.

### Usage Notes

The `SDO_DECODE ( )` function is called once for each dimension to be decoded.

Example 9–5 shows the `SDO_DECODE ( )` function.

#### **Example 9–5**

```
SQL> SELECT
2> SDO_TO_BVALUE(SDO_DECODE(DATA_PT,1),1,6),
3> SDO_TO_BVALUE(SDO_DECODE(DATA_PT,2),-100,100),
4> SDO_TO_DATE(SDO_DECODE(DATA_PT,3))
5> FROM sourcetable1 WHERE SAMPLENAME='SAMPLE1';
```

### Related Topics

- `SDO_TO_BVALUE()` function
- `SDO_TO_DATE()` function



---

## SDO\_ENCODE

### Purpose

This function combines dimensions to create the HHCODE that describes an area or point.

### Syntax

```
SDO_ENCODE (dimension1[,dimension2 ...])
```

### Keywords and Parameters

*dimension* Specifies an expression created by the SDO\_BVALUETODIM or SDO\_DATETODIM functions.  
Data type is RAW.

### Returns

This function returns an HHCODE. The data type is RAW.

### Usage Notes

Consider the following when using this function:

- When encoding dimensions, the order of the dimensions in the parameter list must be consistent for all rows within the table.
- This function can encode up to 32 dimensions.

Example 9–6 shows the SDO\_ENCODE ( ) function.

#### **Example 9–6**

```
SQL> INSERT INTO sourcetable1(SAMPLENAME,DATA_PT)
2> VALUES ('SAMPLE1',SDO_ENCODE(SDO_BVALUETODIM(50,-100,100,10),
3> SDO_BVALUETODIM(30,-100,100,10),
4> SDO_DATETODIM(TO_DATE('05-Jul-96'),3)));
```

### Related Topics

- SDO\_BVALUETODIM() function
- SDO\_DATETODIM() function

## SDO\_TO\_BVALUE

---

### Purpose

This function returns the original bounded data value of a dimension.

### Syntax

SDO\_TO\_BVALUE (*dimension*, *lower\_boundary*, *upper\_boundary*)

### Keywords and Parameters

<i>dimension</i>	Specifies the dimension. Data type is RAW.
<i>lower_boundary</i>	Specifies the lower boundary of the dimension range. Data type is NUMBER.
<i>upper_boundary</i>	Specifies the upper boundary of the dimension range. Data type is NUMBER.

### Returns

This function returns a bounded data value. The data type is NUMBER.

### Usage Notes

This function returns a number that is the value for a dimension within the specified range. This is not necessarily the range for which the dimension was originally created.

Example 9–7 shows the SDO\_TO\_BVALUE ( ) function.

#### **Example 9–7**

```
SQL> SELECT (SDO_TO_BVALUE(SDO_DECODE(DATA_PT,2),-100,100)
2> FROM sourcetable1 WHERE SAMPLENAME='SAMPLE1');
```

### Related Topics

- SDO\_DECODE() function
- SDO\_BVALUETODIM() function

## SDO\_TO\_DATE

### Purpose

This function returns the original date value of a dimension.

### Syntax

SDO\_TO\_DATE (*dimension*)

### Keywords and Parameters

*dimension* Specifies the dimension.  
Data type is RAW.

### Returns

This function returns an Oracle DATE data type.

### Usage Notes

Example 9–8 shows the SDO\_TO\_DATE ( ) function.

#### **Example 9–8**

```
SQL> SELECT SDO_TO_DATE(SDO_DECODE(DATA_PT,3))  
2> FROM sourcetable1 WHERE SAMPLENAME='SAMPLE1';
```

### Related Topics

- SDO\_DATETODIM() function
- SDO\_DECODE() function



---

# Sample SQL Scripts and Tuning Tips

This appendix provides supplemental information to aid in setup, maintenance, and tuning of a spatial database. The scripts and tuning suggestions provided are intended as guidelines that can be adapted to the specific needs of your database.

## A.1 Sample SQL Scripts

Spatial Cartridge provides sample SQL script files to show how to use dynamic SQL in a PL/SQL block to create layer tables for spatially indexed data or to administer and manipulate all the partitions of a partitioned spatial table. The scripts are available after installation in the `ORACLE_HOME/md/admin` directory.

### A.1.1 Scripts for Spatial Indexing

This section describes the `cr_spatial_index.sql` and `crlayer.sql` scripts.

#### A.1.1.1 `cr_spatial_index.sql` Script

The `cr_spatial_index.sql` script file shows an example of updating the spatial index for a layer, and executing a commit after every 50 GIDs have been entered.

The procedures `SDO_ADMIN.POPULATE_INDEX()` and `SDO_ADMIN.POPULATE_INDEX_FIXED()` operate as a single transaction. To reduce the amount of rollback required to execute these procedures, you can write a routine similar to that in `cr_spatial_index.sql`. This script loops and calls `SDO_ADMIN.UPDATE_INDEX_FIXED()` for each GID, committing after every 50 GIDs.

```
-- cr_spatial_index.sql
--
-- Note: if geometries do not span more than 1 row, you can remove
-- the DISTINCT qualifier from the SELECT statement
```

```
--
declare
  cursor c1 is SELECT DISTINCT sdo_gid from POLYGON_SDOGEOM;
  gid number;
  i number;
begin
  i := 0;
  for r in c1 loop
    begin
      gid:= r.sdo_gid;
      sdo_admin.update_index_fixed('POLYGON', gid, 15, FALSE, FALSE, FALSE);
      exception when others then
        dbms_output.put_line('error for gid'||to_char(gid)||': '||SQLERRM );
      end;
      i:= i + 1;
      if i = 50 then
        commit;
        i:= 0;
      end if;
    endloop;
  commit;
end;
/
```

When you call the `SDO_ADMIN.UPDATE_INDEX_FIXED( )` procedure for a large data set, you may get a "snapshot too old" error message from the Oracle server. You can avoid this error by creating more or larger rollback segments. You can also try to increase the number of GIDs before committing the transaction.

---

---

**Note:** The `cr_spatial_index.sql` script is not available in your `ORACLE_HOME/md/admin` directory after installation. You must create this script yourself.

---

---

### A.1.1.2 `crlayer.sql` Script

The `crlayer.sql` script file is a template used to create all the tables for a layer and populate the metadata in the `<layername>_SDODIM` and `<layername>_SDOLAYER` tables.

## A.1.2 Scripts for Partitioned Point Data

This section describes the following scripts:

- `altpart.sql`
- `drppart.sql`
- `sdogrant.sql`

Although the scripts described in this section are available, the recommended approach is to use Oracle8 Partitioning and spatial indexing.

#### A.1.2.1 `altpart.sql` Script

The `altpart.sql` script file shows how to use dynamic SQL in a PL/SQL procedure to modify all partitions of a Spatial Cartridge partitioned table.

The Spatial Cartridge data dictionary view used in this SQL script requires that a registered Spatial Cartridge partitioned table is specified. If the table is not registered, you can use the `USER_TABLES` view to select all the partition tables from the user's schema. To use the `USER_TABLES` view, use the following syntax:

```
SQL> SELECT TABLENAME FROM user_tables WHERE TABLENAME LIKE  
2> '%tablename_P%';
```

#### A.1.2.2 `drppart.sql` Script

The `drppart.sql` script file shows how to use dynamic SQL in a PL/SQL procedure to drop all partitions of a Spatial Cartridge partitioned table. After running this procedure, you must run the `SDO_ADMIN.DROP_PARTITION_INFO()` procedure.

The Spatial Cartridge data dictionary view used in this SQL script requires that a registered Spatial Cartridge partitioned table is specified. If the table is not registered, you can use the `USER_TABLES` view to select all the partition tables from the user's schema. To use the `USER_TABLES` view, use the following syntax:

```
SQL> SELECT TABLENAME FROM user_tables WHERE TABLENAME LIKE  
2> '%tablename_P%';
```

#### A.1.2.3 `sdogrant.sql` Script

The `sdogrant.sql` script file contains an administrative procedure, `PROPAGATE_GRANTS()`, which is used after calling the `SDO_ADMIN.PARTITION()` or `SDO_ADMIN.REPARTITION()` procedures.

This procedure must first be compiled by running the `sdogrant.sql` file. The `PROPAGATE_GRANTS()` procedure is only callable by the user who compiled it.

## A.2 Tuning Tips

The following information can be used as a guideline for tuning a spatial database.

### A.2.1 Data Modeling

Data modeling is very important when designing a spatial database. You should group geometries into layers based on the similarity of their attributes. Assume your data model uses line strings to represent both roads and rivers. The attributes of a road and the attributes of a river are different. Therefore, these geometries should be modeled in two different layers.

In practice, however, if the user of your application will always ask to see both the roads and rivers in a particular region (area of interest), then it may be appropriate to model roads and rivers in the same layer with a common set of attributes.

It is equally important to understand how the data in the various layers will be queried. If the user of your application is interested in querying the data based on a relationship between the layers, then you should index the layers with the same tiling level. For example, a query such as, "which roads cross rivers?" can achieve better performance if the roads and rivers layers are tiled at the same level.

It is not always critical to tile all your layers to the same level. You may find times when you have two layers that are optimally tiled to different levels (for example zipcode boundaries may be tiled to level 5 and major roads may be tiled to level 7). If you want to ask the question, give me all the major roads that intersect a particular zipcode boundary, it is not necessary to retiling all the zipcode boundaries to level 7. You can move the zipcode boundary of interest to a temporary layer and just retiling that one geometry to level 7. You can then perform the query.

### A.2.2 Understanding the Tiling Level

The following example explains how tiling is used in Spatial Cartridge.

Assume you want all the roads (line strings) that overlap a county boundary (polygon) in a spatial database containing 10 million roads. Ignoring Spatial Cartridge features for a moment, in purely mathematical terms, the problem translates into comparing all the line segments that make up each road, to the line segments and area of the county boundary to see if there is any intersection. This geometry-to-geometry comparison is very expensive.

Spatial Cartridge simplifies this calculation by approximating each geometry with fixed-sized tiles. The primary filter in Spatial Cartridge translates the problem to show all the roads that have a tile equal to a tile that approximates the polygon. The result of this is a superset of the final answer.



The secondary filter (a true geometry-to-geometry comparison) can now be applied to the candidates that returned from the Spatial Cartridge primary filter, instead of to every road in the database.

Picking the correct tile size for fixed tiling is one of the most important factors in attaining good performance. If the tile size you select is too small, you could end up generating thousands of tiles per geometry. Also, the process of tiling a query window (like the county boundary mentioned previously) may become very time consuming.

At the same time, you do not want to choose tiles that are too big. This would defeat the purpose of the Spatial Cartridge primary filter. If the tiles are too big, then too many geometries are returned from the primary filter and are sent to the more costly secondary filter.

Keep in mind that the tile size you choose should also depend on whether the query window (area of interest) is already defined in the database. If the query window is defined in the database, (that is, if the spatial tables and spatial indexes already exist), then you should choose a smaller tile size. Assume the State layer and the Highway layer are already defined in the database. You could perform a spatial join query such as, "which interstate highways go through the state?" without incurring the overhead of tiling because the query window is already defined in the database. If, on the other hand, you are creating the query window dynamically, you have to factor in the time it takes to define and index the query window. In this case, you should choose a larger tile size to reduce the time it takes to define and index the query window.

Oracle recommends running the `SDO_TUNE.ESTIMATE_TILING_LEVEL()` function on your data set to get an initial tiling level estimate. This may not be your final answer, but it will be a good level to start your analysis. In general, it is also recommended that you take a random sample of your data and check the query performance at different levels of tiling. This would give an indication of what is the best tiling level for the total data set.

### A.2.3 Database Sizing

Properly choosing rollback segments and tablespaces are important for getting good performance from Spatial Cartridge. Therefore, it is very important to read the *Oracle8 Administrator's Guide* and understand the concepts of tablespaces and rollbacks.

Here are some general guidelines to consider:

- Always make sure that you have enough rollback space to run the `SDO_ADMIN.POPULATE_INDEX_FIXED()` procedure. To reduce rollback seg-

ment size, run the `SDO_ADMIN.UPDATE_INDEX_FIXED()` procedure in a loop (see Section A.1.1.1).

- Create separate tablespaces for Data Layers, Indexes, and Rollback Segments.
- Properly define initial extents, next extents, and `pctincrease` for data layer tables.
- Define the initial extent as small as possible when you create the `<layer-name>_SDOLAYER` and the `<layername>_SDODIM` tables. These tables contain a few rows each and a small initial extent will reduce the amount of wasted space.
- Use the `SDO_ADMIN.VALIDATE_GEOMETRY()` procedure to ensure correctness of geometries in the data sets. Entering incorrect data may lead to unexpected behavior in index creation and in the `SDO_GEOM.RELATE()` functions.
- Always build a B-tree index on the `SDO_GID` column of the `<layer-name>_SDOGEOM` table before attempting to call the `SDO_ADMIN.POPULATE_INDEX_FIXED()`, `SDO_ADMIN.UPDATE_INDEX_FIXED()`, `SDO_ADMIN.POPULATE_INDEX()`, or `SDO_ADMIN.UPDATE_INDEX()` procedure.
- For fixed-size tiling, always build a B-tree index on the `SDO_CODE` column of the `<layername>_SDOINDEX` table before trying any queries using this table.
- Always build a B-tree index on the `SDO_GID` column of the `<layername>_SDOINDEX` table if individual `SDO_GIDs` will be used as query windows for other Spatial Cartridge layers.
- Visualizing the indexing tiles, as described in Section A.2.8, can lead to a greater understanding of the tuning process with respect to the size of the tiles.
- For variable-sized tiling, always build a B-tree index on the `SDO_GROUPCODE` column of the `<layername>_SDOINDEX` table before trying any queries using this table.

## A.2.4 Tuning Point Data

Point data, unlike line and polygon data, has the unique characteristic of containing one tile per point. This section describes how to improve the performance of queries on point data.

### A.2.4.1 Efficient Queries for Point Data

When querying point data with a rectangular query window, you can take advantage of the nature of these geometries to improve performance.

A rectangle can be defined by its lower-left and upper-right coordinates (Xmin, Ymin and Xmax, Ymax). A point has a single set of coordinates (Px, Py). When your area-of-interest is a rectangle, instead of using the `SDO_GEOM.RELATE()` function in the secondary filter, you can use simple SQL comparison operators as follows:

```
SELECT sdo_gid, sdo_x1, sdo_y1
FROM cities_sdogeom,
     (SELECT a.sdo_gid gid1
      FROM cities_sdoindex a,
           window_sdoindex b
      WHERE b.sdo_gid = [area of interest id]
            AND a.sdo_code = b.sdo_code)
     WHERE sdo_gid = gid1
           AND sdo_x1 BETWEEN Xmin AND Xmax
           AND sdo_y1 BETWEEN Ymin AND Ymax;
```

The `DISTINCT` clause is not necessary in the primary filter of the query because a point contains only a single tile in the spatial index.

#### A.2.4.2 Efficient Schema for Point Layers

Because a point contains only one tile in a spatial index, you can place the columns normally found in the `<layername>_SDOINDEX` table in the `<layername>_SDOGEOM` table. This will save you the cost of joining the `<layername>_SDOINDEX` and `<layername>_SDOGEOM` tables. You still need to create an updatable view for the `<layername>_SDOINDEX` table that selects the appropriate columns from the `<layername>_SDOGEOM` table. This is because functions such as `SDO_ADMIN.UPDATE_INDEX_FIXED()` and `SDO_ADMIN.POPULATE_INDEX_FIXED()` expect a `<layername>_SDOINDEX` table to exist. The following example shows a window query of a layer containing point data when the window layer contains one rectangle:

```
SELECT sdo_gid, sdo_x1, sdo_y1
FROM cities_sdogeom a,
     window_sdoindex b
WHERE b.sdo_gid = [area of interest id]
      AND a.sdo_code = b.sdo_code)
      AND sdo_x1 BETWEEN Xmin AND Xmax
      AND sdo_y1 BETWEEN Ymin AND Ymax;
```

## A.2.5 Tuning Spatial Join Queries

There are some helpful hints you can place in your spatial join queries to improve performance. The remainder of this section describes some of the hints you can use. For more information on hints, see the *Oracle8 Tuning* manual.

### A.2.5.1 Using the NO\_MERGE, INDEX, and USE\_NL Hints

A spatial join takes place between two layers. When the two layers being joined are line or polygon layers, the spatial join query contains two DISTINCT clauses: one in the inner SELECT clause and the other in the outer SELECT clause. The Oracle optimizer ignores the inner DISTINCT clause to save on the cost of sorting. However, if the inner DISTINCT clause is ignored, the secondary filter gets called many more times than it needs to be. This can have a significant impact on performance because the secondary filter is an expensive operation. Use the NO\_MERGE hint to prevent the optimizer from ignoring the inner DISTINCT clause.

In a spatial join, all the tiles from one layer are compared to all the tiles from another layer. The Oracle server performs a full table scan on one <layer-name>\_SDOINDEX table, (preferably the smaller of the two), and an index lookup on the other <layername>\_SDOINDEX table. Use the INDEX and USE\_NL hints to force the optimizer to perform the full table scan on the smaller of the two <layer-name>\_SDOINDEX tables being compared.

The following example shows a spatial join between line (road) and polygon (county) data. The query answers the question, "which counties intersect major roads?"

```
SELECT /*+ cost
        ordered use_nl(COUNTY_sdogeom)
        index (COUNTY_sdogeom NAME_OF_SDO_GID_INDEX)
*/
COUNTY_sdogeom.SDO_GID,
COUNTY_sdogeom.SDO_ESEQ,
COUNTY_sdogeom.SDO_SEQ,
COUNTY_sdogeom.SDO_X1,COUNTY_sdogeom.SDO_Y1,
COUNTY_sdogeom.SDO_X2,COUNTY_sdogeom.SDO_Y2,
COUNTY_sdogeom.SDO_X3,COUNTY_sdogeom.SDO_Y3,
COUNTY_sdogeom.SDO_X4,COUNTY_sdogeom.SDO_Y4,
COUNTY_sdogeom.SDO_X5,COUNTY_sdogeom.SDO_Y5,
COUNTY_sdogeom.SDO_X6,COUNTY_sdogeom.SDO_Y6,
COUNTY_sdogeom.SDO_X7,COUNTY_sdogeom.SDO_Y7,
COUNTY_sdogeom.SDO_X8,COUNTY_sdogeom.SDO_Y8
FROM (SELECT DISTINCT gid_a gid1
```

```

FROM (SELECT /*+ index (a NAME_OF_SDO_CODE_INDEX)
            index (b NAME_OF_SDO_CODE_INDEX)
            use_nl (a b)
            no_merge */
      DISTINCT a.sdo_gid gid_a,
              b.sdo_gid gid_b
      FROM COUNTY_SDOINDEX a,
           MAJOR_ROAD_SDOINDEX b
      WHERE a.sdo_code = b.sdo_code)
WHERE sdo_geom.relate('COUNTY', gid_a, 'ANYINTERACT',
                    'MAJOR_ROAD',gid_b) <> 'FALSE'),
COUNTY_sdogeom
WHERE COUNTY_sdogeom.sdo_gid = gid1;

```

### A.2.5.2 Spatial Join Queries with Point Layers

The following example shows a spatial join between line (road) and point (street address) data. The query answers the question, "which addresses are on a major road?"

```

SELECT /*+ cost
        ordered use_nl (STREET_ADDRESS_sdogeom)
        index (STREET_ADDRESS_sdogeom NAME_OF_SDO_GID_INDEX)
        */
      STREET_ADDRESS_sdogeom.SDO_GID,
      STREET_ADDRESS_sdogeom.SDO_X1,
      STREET_ADDRESS_sdogeom.SDO_Y1
FROM (SELECT DISTINCT gid_a gid1
      FROM (SELECT /*+ index (a NAME_OF_SDO_CODE_INDEX)
                    index (b NAME_OF_SDO_CODE_INDEX)
                    use_nl (a b) */
            a.sdo_gid gid_a,
            b.sdo_gid gid_b
            FROM STREET_ADDRESS_SDOINDEX a,
                 MAJOR_ROAD_SDOINDEX b
            WHERE a.sdo_code = b.sdo_code)
      WHERE sdo_geom.relate('STREET_ADDRESS', gid_a, 'ANYINTERACT',
                          'MAJOR_ROAD',gid_b) <> 'FALSE'),
COUNTY_sdogeom
WHERE COUNTY_sdogeom.sdo_gid = gid1;

```

The inner DISTINCT clause is not necessary for spatial joins where one of the layers contains point data. Therefore, the NO\_MERGE hint is not necessary. This is because points contain only one tile in the spatial index.

The following example shows a spatial join between polygon (county) and point (street address) data. The query generates a report that displays how many addresses are associated with each county.

If you can assume that each street address is associated with a single county, you can significantly speed up this query. Because points contain only a single tile in the spatial index, any street address tile that matches only one county tile in the primary filter does not need to go through the expensive secondary filter.

```
SELECT county_gid, count(street_gid)
FROM (SELECT poly.sdo_gid county_gid, street.sdo_gid street_gid
      FROM STREET_ADDRESS_sdoindex street,
           (SELECT sdo_code county_sdo_code,
                count(sdo_gid) interacts
           FROM CENSUS_COUNTY_sdoindex
           GROUP by sdo_code
           ) counts,
      CENSUS_COUNTY_sdoindex poly
 WHERE street.sdo_code = counts.county_sdo_code
       AND poly.sdo_code = street.sdo_code
       AND (counts.interacts = 1
           OR
           sdo_geom.relate('STREET_ADDRESS', street.sdo_gid,
                          'ANYINTERACT',
                          'CENSUS_COUNTY',poly.sdo_gid) <> 'FALSE'
           )
      )
GROUP BY county_gid;
```

## A.2.6 Using Customized Geometry Types

Spatial Cartridge supports three geometry types: points, lines, and polygons. If your data contains another type, such as a circle or arc, then you must choose the supported type that best approximates your desired type. For example, a circle can be defined as a multi-sided polygon. Obviously, the more coordinates in the element, the better the approximation will be.

Although customized types are not supported, you do not have to lose your knowledge of the type. After storing the approximated element, create another element in that geometry with ETYPE=0. Spatial Cartridge ignores elements of ETYPE 0. You can then write your own routines to handle your specialized geometry type.

## A.2.7 Performing Secondary Filter Queries and the Redo Log

When the Oracle server processes SQL statements that require sorting, such as statements containing an ORDER BY or DISTINCT clause, the Oracle server stores the result set in a temporary storage area. The result set is then sorted. If the SORT\_AREA\_SIZE is insufficient for holding the result set in memory, then some data may be written to disk and an entry is written in the redo log.

The RELATE() and INTERACT() secondary filters issue SQL statements internally that contain DISTINCT and ORDER BY clauses. If the SORT\_AREA\_SIZE initialization parameter is too small for processing the secondary filters, then some sorting may occur on disk, which causes entries to be written in the redo log. This may affect performance. For better performance, increase the SORT\_AREA\_SIZE parameter to force sorts to occur in memory.

## A.2.8 Visualizing the Spatial Index (Drawing Tiles)

To select an appropriate tiling level, it may help to visualize the tiles covering your geometries. Through visualization, you can determine how many tiles are used for each object, the size of the tiles, and how well the edges of your geometry are covered. The basic algorithm is:

1. select the edges of the tiles represented by the index entries
2. plot the tiles on a two-dimensional grid
3. plot your geometries on the same grid

The Spatial Cartridge spatial index is represented internally as a linear quadtree. The structure used to represent the linear quadtree is composed of two components: a data component and a metadata component. The data component of the linear quadtree is stored in the SDO\_CODE column, and the metadata component is stored in the SDO\_META column.

The SDO\_META column is not required for spatial queries. However, by combining the SDO\_META column with the SDO\_CODE column, the tiles of any geometry or for the entire data set can be decoded. This capability allows the tiles to be visualized.

Two Spatial Cartridge internal functions have been made visible in order to describe the tiles. These functions were part of a previous release of Oracle Spatial Data Option, and are currently reserved for internal use only. The functions are not recommended for general use, except for this visualization example. Use the following syntax for the internal functions:

```
hhcellbndry(sdo_code || sdo_meta, sdo_dimnum, sdo_lb, sdo_ub,
```

```
hhlength(sdo_code || sdo_meta) {'MIN' | 'MAX'})
```

In the following examples, the dimension boundaries were assumed to be -180 to 180, and -90 and 90. The dimensional information is stored in the <layer-name>\_SDODIM table.

If you used `SDO_ADMIN.UPDATE_INDEX_FIXED()` or `SDO_ADMIN.POPULATE_INDEX_FIXED()` to generate your spatial index, replace "sdo\_code || sdo\_meta" with `sdo_tile` in the SQL statements below.

The following SQL query can be used to decode all the index entries in a <layername>\_SDOINDEX table. The example returns the coordinates of the lower-left and upper-right corners of each tile.

```
SELECT hcellbndry (sdo_code || sdo_meta, 1, -180.000000000, 180.000000000,
    hhlength (sdo_code || sdo_meta), 'MIN') min_x,
    hcellbndry (sdo_code || sdo_meta, 1, -180.000000000, 180.000000000,
    hhlength (sdo_code || sdo_meta), 'MAX') max_x,
    hcellbndry (sdo_code || sdo_meta, 2, -90.000000000, 90.000000000,
    hhlength (sdo_code || sdo_meta), 'MIN') min_y,
    hcellbndry (sdo_code || sdo_meta, 2, -90.000000000, 90.000000000,
    hhlength (sdo_code || sdo_meta), 'MAX') max_y
FROM (SELECT DISTINCT sdo_code, sdo_meta FROM <layer_name>_sdoindex);
```

The following SQL query can be used to decode the index entries for a specific geometry stored in a <layername>\_SDOINDEX table:

```
SELECT hcellbndry (sdo_code || sdo_meta, 1, -180.000000000, 180.000000000,
    hhlength (sdo_code || sdo_meta), 'MIN') min_x,
    hcellbndry (sdo_code || sdo_meta, 1, -180.000000000, 180.000000000,
    hhlength (sdo_code || sdo_meta), 'MAX') max_x,
    hcellbndry (sdo_code || sdo_meta, 2, -90.000000000, 90.000000000,
    hhlength (sdo_code || sdo_meta), 'MIN') min_y,
    hcellbndry (sdo_code || sdo_meta, 2, -90.000000000, 90.000000000,
    hhlength (sdo_code || sdo_meta), 'MAX') max_y
FROM <layer_name>_sdoindex
WHERE sdo_gid = <geometry id>;
```



---

---

## Data Dictionary

The Spatial Cartridge data dictionary is a set of tables owned by the database user `mdsys`. An extension to the Oracle8 data dictionary, it automatically maintains information about spatial tables, columns, and partitions. The Spatial Cartridge data dictionary is created during the installation process. All nonspatial attribute information is maintained in the Oracle8 data dictionary.

The Spatial Cartridge data dictionary has public views that provide extensive information about spatial tables. This appendix contains descriptions of the views that are available.

---

---

**Note:** Only the partitioned point routines use the Spatial Cartridge data dictionary.

---

---

The following views are publicly available:

- `ALL_MD_COLUMNS`
- `ALL_MD_DIMENSIONS`
- `ALL_MD_EXCEPTIONS`
- `ALL_MD_LOADER_ERRORS`
- `ALL_MD_PARTITIONS`
- `ALL_MD_TABLES`
- `ALL_MD_TABLESPACES`
- `DBA_MD_COLUMNS`

- DBA\_MD\_DIMENSIONS
- DBA\_MD\_EXCEPTIONS
- DBA\_MD\_LOADER\_ERRORS
- DBA\_MD\_PARTITIONS
- DBA\_MD\_TABLES
- DBA\_MD\_TABLESPACES
- USER\_MD\_COLUMNS
- USER\_MD\_DIMENSIONS
- USER\_MD\_EXCEPTIONS
- USER\_MD\_LOADER\_ERRORS
- USER\_MD\_PARTITIONS
- USER\_MD\_TABLES
- USER\_MD\_TABLESPACES

---



---

**WARNING:** Do not delete or modify any of the tables in the mdsys account. This corrupts the Spatial Cartridge data dictionary.

---



---

## ALL\_MD\_COLUMNS

Returns a list of all columns that are part of spatial tables.

Column	Description
OWNER	owner of the object
MD_TABLE_NAME	name of the spatial table
COLUMN_NAME	name of the column
DATA_TYPE	data type of the column
DATA_LENGTH	length of the column in bytes
DATA_PRECISION	scale for NUMBER data type, binary precision for FLOAT data type, and NULL for all other data types
DATA_SCALE	digits to right of decimal point in an HHCODE or a number

---

<b>Column</b>	<b>Description</b>
NDIM	number of dimensions in the HHCODE column (It is NULL for all other data types.)
MAX_LEVEL	maximum number of levels in the column
NULLABLE	indicates if column allows NULL values
PARTITION_KEY	indicates if column is the partition key column; only one is allowed per partitioned table
COLUMN_ID	sequence number of the column as created
DEFAULT_LENGTH	length of the default value for the column
NUM_DISTINCT	number of distinct values in each column of the table
LOW_VALUE	lowest value for tables with three or fewer rows (It is the second-lowest value in the column for tables with more than three rows.)
HIGH_VALUE	highest value for tables with three or fewer rows (It is the second-highest value in the column for tables with more than three rows.)

---

### **ALL\_MD\_DIMENSIONS**

Returns a list of all dimensions that are part of HHCODE columns.

<b>Column</b>	<b>Description</b>
OWNER	owner of the object
MD_TABLE_NAME	name of the spatial table
COLUMN_NAME	name of the column
DIMENSION_NAME	name of the dimension
DIMENSION_NUMBER	dimension number
LOWER_BOUND	lower boundary of the dimension range
UPPER_BOUND	upper boundary of the dimension range
SCALE	scale of the dimension
RECURSION_LEVEL	number of levels encoded in the HHCODE

---

---

## ALL\_MD\_EXCEPTIONS

Contains information about spatial tables that should be dropped as a result of some failed operation, such as a failed load.

Column	Description
OWNER	owner of the object
NAME	object name
OPERATION	operation during which the failure occurred
CCHH	common code HHCODE

## ALL\_MD\_LOADER\_ERRORS

Contains the current status of a file that was loaded into a table using SD\*Loader.

Column	Description
OWNER	owner of the object
MD_TABLE_NAME	spatial table name
FILENAME	SLF file name
ROWS_LOADED	number of rows loaded before failure

## ALL\_MD\_PARTITIONS

Returns a list of all the partitioned tables that are part of a user-accessible spatial table.

Column	Description
OWNER	owner of the object
MD_TABLE_NAME	name of the spatial table
PARTITION_TABLE_NAME	name of the partitioned table
CLASS	class of partition: NODE or LEAF
COMMON_LEVEL	number of levels of resolution of the common HHCODE for the partition
COMMON_HHCODE	common HHCODE substring for the partition
OFFLINE_STATUS	status of partition: ONLINE or OFFLINE

---

Column	Description
ARCHIVE_DATE	date of last archive

### ALL\_MD\_TABLES

Returns a list of all the user-accessible spatial tables.

Column	Description
OWNER	owner of the table
MD_TABLE_NAME	name of the spatial table
CLASS	class of table: PARTITIONED or NON-PARTITIONED
PTAB_SEQ	number of last partitioned table created
HIGH_WATER_MARK	maximum number of rows that can be inserted into a partitioned table
OFFLINE_PATH	complete path name to directory where the table is archived
COUNT_MODE	count mode for estimating number of rows in a partition: ESTIMATE or EXACT

### ALL\_MD\_TABLESPACES

Returns a list of all tablespaces used by spatial tables.

Column	Description
OWNER	owner of the object
MD_TABLE_NAME	name of the spatial table
TABLESPACE_NAME	name of tablespace
SEQUENCE	sequence number
STATUS	status of tablespace: ACTIVE or INACTIVE

### DBA\_MD\_COLUMNS

Returns a list of all columns that are part of Spatial Cartridge tables.

Column	Description
OWNER	owner of the object

---

<b>Column</b>	<b>Description</b>
MD_TABLE_NAME	name of the spatial table
COLUMN_NAME	name of the column
DATA_TYPE	data type of the column
DATA_LENGTH	length of the column in bytes
DATA_PRECISION	scale for NUMBER data type, binary precision for FLOAT data type, and NULL for all other data types
DATA_SCALE	digits to right of decimal point in an HHCODE or a number
NDIM	number of dimensions in the HHCODE column (It is NULL for all other data types.)
MAX_LEVEL	maximum number of levels in the column
NULLABLE	indicates if column allows NULL values
PARTITION_KEY	indicates if column is the partition key column; only one is allowed per partitioned table
COLUMN_ID	sequence number of the column as created
DEFAULT_LENGTH	length of the default value for the column
NUM_DISTINCT	number of distinct values in each column of the table
LOW_VALUE	lowest value for tables with three or fewer rows (It is the second-lowest value in the column for tables with more than three rows.)
HIGH_VALUE	highest value for tables with three or fewer rows (It is the second-highest value in the column for tables with more than three rows.)

---

## **DBA\_MD\_DIMENSIONS**

Returns a list of all dimensions that are a part of spatial tables.

<b>Column</b>	<b>Description</b>
OWNER	owner of the object
MD_TABLE_NAME	name of the spatial table
COLUMN_NAME	name of the column
DIMENSION_NAME	name of the dimension

---

<b>Column</b>	<b>Description</b>
DIMENSION_NUMBER	dimension number
LOWER_BOUND	lower boundary of the dimension range
UPPER_BOUND	upper boundary of the dimension range
SCALE	scale of the dimension
RECURSION_LEVEL	number of levels encoded in the HHCODE

### **DBA\_MD\_EXCEPTIONS**

Contains information about spatial tables that should be dropped as a result of some failed operation, such as a failed load.

<b>Column</b>	<b>Description</b>
OWNER	owner of the object
NAME	object name
OPERATION	operation during which the failure occurred
CCHH	common code HHCODE

### **DBA\_MD\_LOADER\_ERRORS**

Contains the current status of a file that was loaded into a table using SD\*Loader.

<b>Column</b>	<b>Description</b>
OWNER	owner of the table where the error occurred
MD_TABLE_NAME	spatial table name
FILENAME	SLF file name
ROWS_LOADED	number of rows loaded before failure

### **DBA\_MD\_PARTITIONS**

Returns a list of all the partitioned tables.

<b>Column</b>	<b>Description</b>
OWNER	owner of the object
MD_TABLE_NAME	name of the spatial table

---

<b>Column</b>	<b>Description</b>
PARTITION_TABLE_NAME	name of the partitioned table
CLASS	class of partition: NODE or LEAF
COMMON_LEVEL	number of levels of resolution of the common HHCODE for the partition
COMMON_HHCODE	common HHCODE substring for the partition
OFFLINE_STATUS	status of partition: ONLINE or OFFLINE
ARCHIVE_DATE	date of last archive

---

### **DBA\_MD\_TABLES**

Returns a list of all the spatial tables.

<b>Column</b>	<b>Description</b>
OWNER	owner of the table
MD_TABLE_NAME	name of the spatial table
CLASS	class of table: PARTITIONED or NON-PARTITIONED
PTAB_SEQ	number of last partitioned table created
HIGH_WATER_MARK	maximum number of rows that can be inserted into a partitioned table
OFFLINE_PATH	complete path name to directory where the table is archived
COUNT_MODE	count mode for estimating number of rows in a partition: ESTIMATE or EXACT

---

### **DBA\_MD\_TABLESPACES**

Returns a list of all tablespaces used by spatial tables.

<b>Column</b>	<b>Description</b>
OWNER	owner of the object
MD_TABLE_NAME	name of the spatial table
TABLESPACE_NAME	name of tablespace
SEQUENCE	sequence number
STATUS	status of tablespace: ACTIVE or INACTIVE

---



---

## USER\_MD\_COLUMNS

Returns a list of all the HHCODE columns that are part of tables owned by the user.

Column	Description
MD_TABLE_NAME	name of the spatial table
COLUMN_NAME	name of the spatial table
DATA_TYPE	data type of the column
DATA_LENGTH	length of the column in bytes
DATA_PRECISION	scale for NUMBER data type, binary precision for FLOAT data type, and NULL for all other data types
DATA_SCALE	digits to right of the decimal point in an HHCODE or a number
NDIM	number of dimensions in the HHCODE column (It is NULL for all other data types.)
MAX_LEVEL	maximum number of levels in the column
NULLABLE	indicates if column allows NULL values
PARTITION_KEY	indicates if column is the partition key column; only one allowed per partitioned table
COLUMN_ID	sequence number of the column as created
DEFAULT_LENGTH	length of the default value for the column
NUM_DISTINCT	number of distinct values in each column of the table
LOW_VALUE	lowest value for tables with three or fewer rows (It is the second-lowest value in the column for tables with more than three rows.)
HIGH_VALUE	highest value for tables with three or fewer rows (It is the second-highest value in the column for tables with more than three rows.)

## USER\_MD\_DIMENSIONS

Returns a list of all dimensions that are part of HHCODE columns owned by the user.

Column	Description
MD_TABLE_NAME	name of the spatial table

---

<b>Column</b>	<b>Description</b>
COLUMN_NAME	name of the column
DIMENSION_NAME	name of the dimension
DIMENSION_NUMBER	dimension number
LOWER_BOUND	lower boundary of dimension range
UPPER_BOUND	upper boundary of dimension range
SCALE	scale of the dimension
RECURSION_LEVEL	number of levels encoded in the HHCODE

### **USER\_MD\_EXCEPTIONS**

Contains information about spatial tables that should be dropped as a result of some failed operation, such as a failed load.

<b>Column</b>	<b>Description</b>
NAME	object name
OPERATION	operation during which the failure occurred
CCHH	common code HHCODE

### **USER\_MD\_LOADER\_ERRORS**

Contains the current status of a file that was loaded into a table using SD\*Loader.

<b>Column</b>	<b>Description</b>
MD_TABLE_NAME	spatial table name
FILENAME	SLF file name
ROWS_LOADED	number of rows loaded before failure

### **USER\_MD\_PARTITIONS**

Returns a list of all the partitioned tables that are part of spatial tables owned by the user.

<b>Column</b>	<b>Description</b>
MD_TABLE_NAME	name of the spatial table

---

<b>Column</b>	<b>Description</b>
PARTITION_TABLE_NAME	name of the partition
CLASS	class of partition: NODE or LEAF
COMMON_LEVEL	number of levels of resolution of the common HHCODE for the partition
COMMON_HHCODE	common HHCODE substring for the partition
OFFLINE_STATUS	status of partition: ONLINE or OFFLINE
ARCHIVE_DATE	date of last archive

### **USER\_MD\_TABLES**

Returns a list of all the spatial tables owned by the user.

<b>Column</b>	<b>Description</b>
MD_TABLE_NAME	name of the spatial table
CLASS	class of table: PARTITIONED or NON-PARTITIONED
PTAB_SEQ	number of last sequence created
HIGH_WATER_MARK	maximum number of rows that can be inserted into a partitioned table
OFFLINE_PATH	complete path name to directory where the table is archived
COUNT_MODE	count mode for estimating number of rows in a partition: ESTIMATE or EXACT

### **USER\_MD\_TABLESPACES**

Returns a list of all tablespaces used by spatial tables.

<b>Column</b>	<b>Description</b>
MD_TABLE_NAME	name of the spatial table
TABLESPACE_NAME	name of tablespace
SEQUENCE	sequence number
STATUS	status of the tablespace: ACTIVE or INACTIVE



---

## Messages and Codes

**MDSQL-00001: partition is OFFLINE**

**Cause:** An MDSQL operation was attempted on a partition that is OFFLINE.

**Action:** Restore the partition and try the operation again.

**MDSQL-00002: PK is out of bounds**

**Cause:** The partition key for the record being inserted belongs in another partition.

**Action:** Insert the record into the correct partition. The correct partition can be identified using the GET\_PARTITION\_NAME() function.

**MDSQL-00003: updates that move the PK are not supported**

**Cause:** The update of the partition key would result in the record belonging to another partition.

**Action:** Use the MD\_DML.MOVE\_RECORD() procedure to update the partition key and move the record to the correct partition.

**SDO-00000 successful completion**

**Cause:** An operation has completed normally, having met no exceptions.

**Action:** No user action required.

**SDO-00002 too many errors in '%s' - aborting**

**Cause:** The max errors allowed has been reached.

**Action:** Fix errors and retry or increase the max errors allowed.

**SDO-00200 could not allocate memory for initialization of internal heap manager**

**Cause:** Could not allocate the essential memory.

---

**Action:** Increase amount of memory available or wait until more memory becomes available.

**SDO-00201 failed to initialize message handler**

**Cause:** Possibly could not allocate the essential memory.

**Action:** Increase amount of memory available or wait until more memory becomes available.

**SDO-00202 out of memory while performing essential allocations**

**Cause:** Could not allocate the essential memory.

**Action:** Increase amount of memory available or wait until more memory becomes available.

**SDO-00203 invalid maximum bind array size**

**Cause:** The space allocated to do the necessary processing was not enough.

**Action:** Increase the value of the bind size.

**SDO-00250 unable to open '%s' for processing**

**Cause:** Could not open the named file for processing.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00251 cannot close '%s'**

**Cause:** Could not close the named file.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00252 unable to read '%s'**

**Cause:** Could not read the named file for processing.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00253 unable to write to '%s'**

**Cause:** Could not write to the named file.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00254 unable to seek in '%s'**

**Cause:** Could not seek the named file for processing.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00255 unable to write to log file**

**Cause:** Could not write to log file.

---

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00256 unable to delete ‘%s’**

**Cause:** Could not delete the named file.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00257 unable to rename ‘%s’**

**Cause:** Could not rename the named file.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00258 unable to open a temporary file for processing**

**Cause:** Could not open the named file for processing.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00259 cannot close a temporary file**

**Cause:** Could not close the named file.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00260 unable to read from a temporary file**

**Cause:** Could not read the named file for processing.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00261 unable to write to a temporary file**

**Cause:** Could not write to the named file.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00262 unable to seek in a temporary file**

**Cause:** Could not seek the named file for processing.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00263 unable to create a unique temporary file name**

**Cause:** Could not create a unique file name.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00264 unable to delete a temporary file**

**Cause:** Could not delete a temporary file.

**Action:** Check the operating system message(s) accompanying this message.

---

**SDO-00265 unable to rename a temporary file**

**Cause:** Could not rename a temporary file.

**Action:** Check the operating system message(s) accompanying this message.

**SDO-00267 failed to translate environment variable in path**

**Cause:** The environment variable is invalid or not defined.

**Action:** Verify that the environment variable is correct.

**SDO-01000 invalid keyword or value for argument number %s**

**Cause:** Incorrect spelling of keyword, keyword is not valid, or incorrect type for argument's value.

**Action:** Check the command syntax and your spelling, then retry.

**SDO-01001 mandatory keyword is missing**

**Cause:** Keyword was left out on the command line.

**Action:** Place mandatory keyword on command line.

**SDO-01002 argument for the keyword '%s' is invalid**

**Cause:** Argument is either missing or out of bounds.

**Action:** Verify argument is correct for the specified keyword.

**SDO-01003 '%s' is not a valid keyword or missing value for the keyword**

**Cause:** Incorrect spelling of keyword, keyword is not valid, or value of the keyword is missing.

**Action:** Check the command syntax and your spelling, then retry.

**SDO-01004 invalid combination on command line**

**Cause:** A keyword was used that could not be combined with a previous keyword on the command line.

**Action:** Check the command line syntax.

**SDO-01005 error detected on the command line**

**Cause:** A command line parameter is improperly defined.

**Action:** Verify that all command line parameters are properly specified.

**SDO-01006 a userid was not specified on the command line**

**Cause:** A user ID was not specified on the command line.



---

**Action:** A user ID must be specified on the command line (for example: USE-RID=sims/sims.)

**SDO-01007 failed to parse connect string**

**Cause:** The connect string specified could not be properly parsed.

**Action:** Verify that the connect string is valid. If it is, document messages and contact Oracle Worldwide Customer Support Services.

**SDO-01008 a null userid was specified**

**Cause:** The user specified a null user ID.

**Action:** A user ID must be entered.

**SDO-01009 a null password was specified**

**Cause:** The user specified a null password.

**Action:** A password must be entered.

**SDO-01200 internal error: [%s]**

**Cause:** An error occurred during an internal match manipulation.

**Action:** Document messages and contact Oracle Worldwide Customer Support Services.

**SDO-01201 internal error: [%s]**

**Cause:** An error occurred during an internal substring manipulation.

**Action:** Document messages and contact Oracle Worldwide Customer Support Services.

**SDO-01202 internal error: [%s]**

**Cause:** An error occurred during an internal increment manipulation.

**Action:** Document messages and contact Oracle Worldwide Customer Support Services.

**SDO-01203 internal error: [%s]**

**Cause:** An error occurred during an internal set ID manipulation.

**Action:** Document messages and contact Oracle Worldwide Customer Support Services.

---

**SDO-01204 internal error: [%s]**

**Cause:** An error occurred during an internal get ID manipulation.

**Action:** Document messages and contact Oracle Worldwide Customer Support Services.

**SDO-01500 invalid username/password; logon denied**

**Cause:** An invalid user name or password was entered in an attempt to log on to Oracle.

**Action:** Enter a valid user name/password combination in the correct format.

**SDO-01501 failed to log onto Oracle**

**Cause:** Log onto Oracle failed.

**Action:** Verify that the Oracle connect string is valid.

**SDO-01502 failed to connect as MDSYS**

**Cause:** The connection to the Oracle MDSYS account failed.

**Action:** Verify that this account exists and was set up correctly.

**SDO-01503 partition is locked**

**Cause:** The partition is currently being used.

**Action:** Wait until the partition is free.

**SDO-01507 partition is locked with an unknown status**

**Cause:** The status of partition is not a known status.

**Action:** Verify that the Spatial Cartridge data dictionary is correct.

**SDO-01508 inserted a partition record after someone else did**

**Cause:** Multiple concurrent loads created same partition (only one wins).

**Action:** Re-run the load process.

**SDO-01600 failed to create table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to create a table.

**SDO-01601 failed to create temporary table**

**Cause:** Check the Oracle messages accompanying this message.

---

**Action:** Verify that you are able to create a table.

**SDO-01602 failed to drop table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to drop a table.

**SDO-01603 failed to drop a temporary table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to drop a table.

**SDO-01604 failed to truncate table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to truncate a table.

**SDO-01605 failed to truncate temporary table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to truncate a table.

**SDO-01606 real table does not exist**

**Cause:** An attempt was made to archive in an archive file and the truncated table does not exist.

**Action:** Document messages and contact Oracle Worldwide Customer Support Services.

**SDO-01607 HHCODE column has no dimensional information**

**Cause:** The dimensional information for the HHCODE column was not found.

**Action:** Verify that the dimensional information is correct.

**SDO-01608 HHCODE column is missing dimensional information in '%s'**

**Cause:** No dimensional information was entered.

**Action:** Verify that the dimensional information is correct.

**SDO-01609 spatial table name does not exist for specified owner**

**Cause:** Either the spatial table name is misspelled or it does not exist.

**Action:** Verify that table exists and check spelling.

---

**SDO-01610 null was passed in not null field ‘%s’**

**Cause:** The null flag was set for a column that was set as a not null field.

**Action:** Set flag for not null setting.

**SDO-01611 dimensional value for ‘%s’ is out of bounds**

**Cause:** The dimensional value is not in the range of the lower and upper bound of the dimension.

**Action:** Verify that the data is correct or if the lower and upper bound of the dimension is correct.

**SDO-01612 the column ‘%s’ has an unsupported data type**

**Cause:** The column is defined with an unsupported data type.

**Action:** Use another data type for this column.

**SDO-01613 ‘%s’ was not found as a non-partitioned spatial table**

**Cause:** The specified table does not exist or is not a non-partitioned spatial table.

**Action:** Verify that specified table exists and is a non-partitioned spatial table.

**SDO-01614 failed to find dimension information from temporary table**

**Cause:** The table may not exist.

**Action:** Verify that table exists and is working properly.

**SDO-01615 failed to find HHCODE column info. from temporary column table**

**Cause:** The temporary table is incorrect.

**Action:** Verify that table exists and is working properly.

**SDO-01616 the spatial table contains no HHCODE columns**

**Cause:** No HHCODE columns exist in specified spatial table.

**Action:** Verify that the HHCODE column exists in specified spatial table.

**SDO-01617 there is no partition key column for the specified table**

**Cause:** There was no HHCODE column specified as the partition key.

**Action:** Alter the spatial table to have a partition key.

---

**SDO-01618 dimensional information was not found for a HHCODE column**

**Cause:** Dimensional information does not exist for a HHCODE column.

**Action:** Verify the Spatial Cartridge data dictionary for the corresponding spatial table.

**SDO-01619 past the maximum level of subdivision**

**Cause:** The high-water mark is being exceeded on a partition at the maximum subdivision point.

**Action:** Alter the spatial table with a higher high-water mark.

**SDO-01620 failed to assign a transaction to specified rollback segment**

**Cause:** The specified rollback segment is invalid.

**Action:** Verify that the rollback segment is correct.

**SDO-01621 failed to create temporary view**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to drop a VIEW.

**SDO-01622 failed to drop temporary view**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to drop a VIEW.

**SDO-01700 failed to create base trigger on partition**

**Cause:** Failed to create a trigger.

**Action:** Check the Oracle message accompanying this message.

**SDO-01701 failed to enable base trigger on partition**

**Cause:** Failed to enable a trigger.

**Action:** Check the Oracle messages accompanying this message.

**SDO-01702 failed to select from table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to select from a table.

---

**SDO-01703 failed to select from temporary table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to select from a table.

**SDO-01704 failed to select information from view user\_users**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to select from this view.

**SDO-01705 failed to select from temporary column table**

**Cause:** The table may not exist.

**Action:** Verify that the table exists and is working properly.

**SDO-01706 failed to select from temporary table**

**Cause:** The temporary table may not exist.

**Action:** Verify that the temporary table exists and is working properly.

**SDO-01707 failed to insert into temporary table**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that you are able to insert into a table.

**SDO-01708 failed to select information from view all\_tables**

**Cause:** Either VIEW does not exist or you do not have access.

**Action:** Verify that you have access to the VIEW and it is working properly.

**SDO-01800 failed to select from sys.col\$**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that MDSYS can select from sys.col\$.

**SDO-01801 failed to select from sys.obj\$**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that MDSYS can select from sys.obj\$.

**SDO-01802 failed to select from sys.ts\$**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that MDSYS can select from sys.ts\$.

---

**SDO-01803 failed to select from sys.tab\$**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that MDSYS can select from sys.tab\$.

**SDO-01804 failed to select from sys.dba\_tables**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that MDSYS can select from sys.dba\_tables.

**SDO-01805 failed to select information from sys.dba\_users**

**Cause:** Check the Oracle messages accompanying this message.

**Action:** Verify that MDSYS can select from sys.dba\_users.

**SDO-01850 invalid HHCODE column name in '%s'**

**Cause:** The HHCODE column name was not found for specified spatial table.

**Action:** Verify the HHCODE column exist in the specified spatial table.

**SDO-01851 no column information was found in view all\_md\_columns**

**Cause:** No non-HHCODE columns were found in specific spatial table.

**Action:** Verify that the view ALL\_MD\_COLUMNS is correct.

**SDO-01852 no class information was found in view all\_md\_tables**

**Cause:** The class information for the spatial table was not found.

**Action:** Verify that the view ALL\_MD\_TABLES is correct.

**SDO-01853 failed to select information from view all\_md\_tables**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary

**SDO-01854 HHCODE column name '%s' was not found for spatial table**

**Cause:** The HHCODE column name was not found for specified spatial table.

**Action:** Verify the HHCODE column exists in the specified spatial table.

---

**SDO-01855 dimension name ‘%s’ was not found for specified HHCODE column**

**Cause:** Either the dimension name does not exist in the specified spatial table or the name is incorrectly spelled.

**Action:** Verify that the given dimension name is correct or verify that it exists in the specified spatial table for the specified HHCODE column.

**SDO-01856 column name ‘%s’ was not found for specified spatial table**

**Cause:** Either the column name does not exist in the specified spatial table or the name is incorrectly spelled.

**Action:** Verify that the given column name is correct and exists in the specified spatial table.

**SDO-01857 dimension name ‘%s’ was not found for partition key column**

**Cause:** Either the dimension name does not exist in the specified spatial table or the name is incorrectly spelled.

**Action:** Verify that the given dimension name is correct or verify that it exists in the specified spatial table for the specified partition key column.

**SDO-01858 not null field in ‘%s’ not passed**

**Cause:** A column value was not passed that was set as a not null field.

**Action:** Pass this column with a value other than null.

**SDO-01859 data type for ‘%s’ is invalid**

**Cause:** The COL structure data type description does not correspond to the data type describing the spatial table.

**Action:** Verify that the data type describing the data is correct or the column in the spatial table has the correct data type.

**SDO-01860 partition name sequence overflow**

**Cause:** Went through all possible values of the sequence.

**Action:** Document messages and contact Oracle Worldwide Support.

**SDO-01900 failed to select information from table md\$col**

**Cause:** Either the table does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.



---

**SDO-01901 failed to select information from table md\$dim**

**Cause:** Either the table does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01902 failed to select information from table md\$ptab**

**Cause:** Either the table does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01903 failed to select information from view mdv\$tab**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01904 failed to select information from view mdv\$col**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01905 failed to select information from view mdv\$dim**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01906 failed to select information from view all\_md\_dimensions**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01907 failed to select information from view all\_md\_columns**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01908 failed to select information from view all\_tab\_columns**

**Cause:** Either the view does not exist or you do not have access.

**Action:** Verify that you have access to the view and it is working properly.

**SDO-01909 failed to find dimension information from view all\_md\_dimensions**

**Cause:** The view may not exist.

**Action:** Verify that the view exists and is working properly.

---

**SDO-01910 failed to find column information from view all\_md\_columns**

**Cause:** The view may not exist.

**Action:** Verify that view exists and is working properly.

**SDO-01911 failed to select information from table md\$ler**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01912 failed to select information from table md\$pts**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01913 failed to select information from table md\$tab**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01914 failed to insert information from table md\$exc**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01915 failed to insert information from table md\$ler**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01916 failed to insert information from table md\$ptab**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01917 failed to delete information from table md\$exc**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01918 failed to delete information from table md\$ler**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

---

**SDO-01919 failed to delete information from table md\$ptab**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01920 failed to update information in table md\$ptab**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01921 failed to update information from table md\$ler**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01922 failed to update information from table md\$pts**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01923 failed to update information from table md\$tab**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01924 failed to select information from view dba\_md\_columns**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-01925 failed to select information from view dba\_md\_dimensions**

**Cause:** Either the view does not exist or is inconsistent.

**Action:** Verify the Spatial Cartridge data dictionary.

**SDO-05000 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The SLF file is built for either a partitioned spatial table or non-partitioned spatial table. The SLF file is being loaded into the wrong class.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

---

**SDO-05001 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The total number of HHCODE columns in SLF file does not match the total in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05002 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The total number of dimensions for each HHCODE column in the SLF file does not match the total in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05003 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The total number of non-HHCODE columns in the SLF file does not match the total in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05004 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The HHCODE column name in SLF file does not match the name in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05005 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The number of dimensions of HHCODE column in the SLF file does not match the number in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

---

**SDO-05006 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The maximum level of dimensions of HHCODE column in the SLF file does not match the maximum level in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05007 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The size of the HHCODE column in the SLF file does not match size in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05008 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The null field of HHCODE column in SLF file does not match null field in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05009 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The dimension name of the HHCODE column in the SLF file does not match name in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05010 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The dimension number of the HHCODE column in the SLF file does not match number in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

---

**SDO-05011 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The dimension level of the HHCODE column in the SLF file does not match level in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05012 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The dimension lower bound of the HHCODE column in the SLF file does not match the lower bound in spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05013 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The dimension upper bound of the HHCODE column in the SLF file does not match the upper bound in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05014 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The column name in the SLF file does not match name in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05015 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The null field of column in the SLF file does not match the null field in spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

---

**SDO-05016 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The column type in the SLF file does not match the type in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05017 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The precision field of a column in the SLF file does not match the precision field in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Cause:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05018 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The scale field of a column in the SLF file does not match the scale field in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05019 SLF header: %s does not match spatial data dictionary: %s**

**Cause:** The size of a column in the SLF file does not match the size in the spatial table. Either the spatial table definition was changed after the SLF file was created or the SLF file is being loaded into the wrong spatial table.

**Action:** Verify the SLF file is being loaded into the correct spatial table or regenerate the SLF file.

**SDO-05020 no data found in SLF file (just header information)**

**Cause:** During the conversion process the writing SLF record process died.

**Action:** Re-create the SLF file verifying that the SLF data is written.

**SDO-05200 line in '%s' was too complex to parse**

**Cause:** The data in the control file exceeded the maximum field limits.

**Action:** Document messages and contact Oracle Worldwide Customer Support Services.

---

**SDO-05201 insufficient number of fields in ‘%s’**

**Cause:** There were fewer than three fields in the control file.

**Action:** Check the control file for accuracy.

**SDO-05202 did not find the expected line or keyword in ‘%s’**

**Cause:** The expected keyword or line was not found.

**Action:** Verify that the specified keywords are in the correct order.

**SDO-05203 an invalid number of partition key columns was specified in ‘%s’**

**Cause:** More than one partition key was specified in the named control file.

**Action:** Verify only one HHCODE is specified as the partition key in the named control file.

**SDO-05204 HHCODE column is missing a dimension sequence in ‘%s’**

**Cause:** When entering the dimensional information into the named control file, a dimensional sequence was left out.

**Action:** Verify that the sequence values for the dimensions are correct.

**SDO-05205 file type keyword not found in ‘%s’**

**Cause:** The required file type keyword in the named file is missing.

**Action:** Verify that the control file has the correct information.

**SDO-05206 unexpected keyword or value found in ‘%s’**

**Cause:** There was more than one field found in the control file.

**Action:** Check the control file for accuracy.

**SDO-05207 ‘%s’ is an invalid file type specified in control file**

**Cause:** The file type was spelled incorrectly. The only valid file types are BINARY and ASCII.

**Action:** Verify that the file type specified in the control file is correct.

**SDO-05208 format keyword not found in ‘%s’**

**Cause:** The required format keyword in the named file is missing.

**Action:** Verify that the control file has the correct information.



---

**SDO-05209 unexpected keyword or value found in ‘%s’**

**Cause:** There were more than two fields found in the control file.

**Action:** Check the control file for accuracy.

**SDO-05210 ‘%s’ is an invalid keyword specified in control file**

**Cause:** The keyword FIXED was not found in control file.

**Action:** Verify that the keyword specified in the control file is correct.

**SDO-05211 ‘%s’ is an invalid record length specified for the data file**

**Cause:** The record length specified is invalid for the given data file.

**Action:** Verify that the record length specified is correct.

**SDO-05212 ‘%s’ is an invalid keyword specified in control file**

**Cause:** The only valid keywords are DIMENSION and COLUMN.

**Action:** Verify that the file type specified in the control file is correct.

**SDO-05213 ‘%s’ is an invalid keyword specified in control file**

**Cause:** The missing keyword POSITION was not found in the control file.

**Action:** Verify that the keyword specified in the control file is correct.

**SDO-05214 invalid position value specified in ‘%s’**

**Cause:** The field position was incorrectly specified.

**Action:** Verify that the field positions specified are correct.

**SDO-05215 ‘%s’ is an invalid data type specified in control file**

**Cause:** The data type specified was spelled incorrectly. The only valid data types are INTEGER, SMALLINT, FLOAT, DOUBLE, BYTEINT, DATE, RAW, and CHAR.

**Action:** Verify that the data type specified is correct.

**SDO-05216 specified position length is not compatible with ‘%s’**

**Cause:** The field position value was incorrectly specified.

**Action:** Verify that the field positions specified are correct.

---

**SDO-05217 '%s' is an invalid keyword specified in control file**

**Cause:** The missing keyword NULLIF was not found in control file.

**Action:** Verify that the keyword specified in the control file is correct.

**SDO-05218 '%s' is an invalid null indicator in control file**

**Cause:** The null indicator was incorrectly specified.

**Action:** Verify that the null indicator is correct.

**SDO-05219 no spatial information was found in '%s'**

**Cause:** Spatial information was not found in the named file.

**Action:** Enter the dimension information into the named file.

**SDO-05220 no dimensional or column information was found in '%s'**

**Cause:** No information was found in the named file.

**Action:** Enter the column and dimension information into the named file.

**SDO-05221 '%s' is not the expected value in the data file**

**Cause:** The expected value was not found.

**Action:** Verify that the specified file is correct.

**SDO-05222 data conversion with column '%s' failed**

**Cause:** Either bad data was passed or incorrect format string was used.

**Action:** Verify that the data and format string are correct.

**SDO-05500 user aborted SD\*Loader**

**Cause:** The user aborted the program.

**Action:** No action required.

**SDO-05501 SQL\*Loader failed to load all records to the partition**

**Cause:** Some records were bad or rejected.

**Action:** Check the log file and correct any errors indicated.

**SDO-05502 SQL\*Loader failed**

**Cause:** The SQL\*Loader process died.

**Action:** Check log file and correct any errors indicated.

---

**SDO-07510 bad boolean value**

**Cause:** A bad Boolean value was specified.

**Action:** Verify that the value is TRUE or FALSE.

**SDO-07511 failed to read user input from terminal**

**Cause:** An error was encountered while attempting to prompt the user for input from the terminal.

**Action:** This is an internal error. Document all messages and contact Oracle Worldwide Customer Support Services.



---

---

# Glossary

## **area**

An extent or region of dimensional space.

## **attribute**

Descriptive information characterizing a geographical feature such as a point, line, or area.

## **attribute data**

Nondimensional data that provides additional descriptive information about multi-dimensional data, for example a class or feature such as a bridge or a road.

## **boundary**

1. The lower or upper extent of the range of a dimension, expressed by a numeric value.
2. The line representing the outline of a polygon.

## **Cartesian coordinate system**

A coordinate system in which the location of a point in n-dimensional space is defined by distances from the point to the reference plane. Distances are measured parallel to the planes intersecting a given reference plane.

## **contain**

To describe a geometric relationship where one object encompasses another and the inner object does not touch any boundaries of the outer. The outer object *contains* the inner object. *See also* inside.

**coordinate**

A set of values uniquely defining a point in an n-dimensional coordinate system.

**coordinate system**

A reference system for the unique definition for the location of a point in n-dimensional space.

**cover**

To describe a geometric relationship in which one object encompasses another and the inner object touches the boundary of the outer object in one or more places.

**data dictionary**

A repository of information about data. A data dictionary stores relational information on all the objects in a database.

**decompose**

To separate or resolve into constituent parts or elements, or into simpler compounds.

**dimensional data**

Data that has one or more dimensional components and is described by multiple values.

**disjoint**

A geometric relationship where two objects do not interact in any way. Two *disjoint* objects do not share any element or piece of their geometry.

**equal**

A geometric relationship in which two objects are considered to represent the same geometric figure. The two objects must be composed of the same number of points, however, the ordering of the points defining the two objects' geometries may differ (clockwise or counter-clockwise).

**extent**

A rectangle bounding a map, the size of which is determined by the minimum and maximum map coordinates.

**feature**

An object in a spatial database with a distinct set of characteristics.

**geographically referenced data**

*See* spatiotemporal data.

**georeferenced data**

*See* spatiotemporal data.

**geographical information system**

A computerized database management system used for the capture, conversion, storage, retrieval, analysis, and display of spatial data.

**GIS**

*See* geographical information system.

**grid**

A data structure composed of points located at the nodes of an imaginary grid. The spacing of the nodes is constant in both the horizontal and vertical directions.

**HHCODE**

A data type representing the intersection point of multiple dimensions. It encodes these multiple dimensions into a unique, linear value. The HHCODEs are used for both spatial indexing and partitioned point data.

**high-water mark**

Expressed in number of records and associated with Spatial Cartridge partitioned table structure, it defines the maximum number of records to store in a table before decomposing another level. The high water mark determines the maximum size of a partition within the Spatial Cartridge table. Partitioned tables are an alternative to spatial indexing.

**hole**

A polygon can include subelements that negate sections of its interior. For example, consider a polygon representing a map of a buildable land with an inner polygon (a hole) representing where a lake is located.

**homogeneous**

Spatial data of one feature type such as points, lines, or regions.

**hyperspatial data**

In mathematics, any space comprising more than the three standard x, y, and z dimensions, also referred to as multidimensional data.

**index**

Identifier that is not part of a database and used to access stored information.

**inside**

To describe a geometric relationship where one object is surrounded by a larger object and the inner object does not touch the boundary of the outer. The smaller object is *inside* the larger. *See also* contain.

**key**

A field in a database used to obtain access to stored information.

**keyword**

Synonym for reserved word.

**latitude**

North/South position of a point on the Earth defined as the angle between the normal to the Earth's surface at that point and the plane of the equator.

**line**

A geometric object represented by a series of points, or inferred as existing between two coordinate points.

**longitude**

East/West position of a point on the Earth defined as the angle between the plane of a reference meridian and the plane of a meridian passing through an arbitrary point.

**multidimensional data**

*See* hyperspatial data.

**partition**

1. The spatial table that contains data only for a unique bounded n-dimensional space.
2. The process of grouping data into partitions that maintain the dimensional organization of the data.



**partition key column**

The primary HHCODE column that is used to dimensionally partition the data. One HHCODE data type column must be identified as the partition key for the table to be registered as partitionable in the Spatial Cartridge data dictionary. There can only be one partition key per spatial table. Note that this is only used for partitioned point data, and not spatially indexed data.

**partitioned table**

The spatial logical table structure that contains one or more partitions. Use partitioned tables only if you are dealing with a very large amount of point data (over 50 GB).

**polygon**

A class of spatial objects having a nonzero area and perimeter, and representing a closed boundary region of uniform characteristics.

**proximity**

A measure of inter-object distance.

**query**

A set of conditions or questions that form the basis for the retrieval of information from a database.

**query window**

Area within which the retrieval of spatial information and related attributes is performed.

**RDBMS**

*See* Relational Database Management System.

**recursion**

A process, function, or routine that executes continuously until a specified condition is met.

**region**

An extent or area of multidimensional space.

## **Relational Database Management System (RDBMS)**

A computer program designed to store and retrieve shared data. In a relational system, data is stored in tables consisting of one or more rows, each containing the same set of columns. Oracle8 is a relational database management system. Other types of database systems are called hierarchical or network database systems.

### **resolution**

The number of subdivision levels of data.

### **scale**

1. The number of digits to the right of the decimal point in a number representing the level of resolution of an HHCODE.
2. The ratio of the distance on a map, photograph, or image to the corresponding image on the ground, all expressed in the same units.

### **SD\*Converter**

A utility used with previous versions of Spatial Data Option to prepare data for loading into spatial tables. Loading is now accomplished through SQL\*Loader.

### **SQL\*Loader**

A utility used with previous versions of Spatial Data Option to load formatted data into spatial tables. Loading is now accomplished through SQL\*Loader.

### **SLF**

See Spatial Load Format.

### **sort**

The operation of arranging a set of items according to a key that determines the sequence and precedence of items.

### **spatial**

A generic term used to reference the mathematical concept of n-dimensional data.

### **spatial data**

Data that is referenced by its location in n-dimensional space. The position of spatial data is described by multiple values. See also hyperspatial data.

### **spatial database**

A database containing information indexed by location.

**spatial data model**

A model of how objects are located on a spatial context.

**Spatial Cartridge data dictionary**

An extension of the Oracle8 data dictionary. It keeps track of the number of partitions created in a spatial table. The Spatial Cartridge data dictionary is owned by MDSYS. The data dictionary is used only by the partitioned point routines.

**spatial data structures**

A class of data structures designed to store spatial information and facilitate its manipulation.

**Spatial Load Format (SLF)**

The format used to load data into spatial tables in a previous version of Spatial Data Option. Loading is now accomplished with the standard SQL\*Loader.

**spatial query**

A query that includes criteria for which selected features must meet location conditions.

**spatiotemporal data**

Data that contains time and/or location components as one of its dimensions, also referred to as geographically referenced data or georeferenced data.

**tessellation**

The process of covering a geometry with rectangular tiles without gaps or overlaps.

**tiling**

*See* tessellation.

**touch**

To describe a geometric relationship where two objects share a common point on their boundaries, but their interiors do not intersect.



## A

---

ADD\_NODES 7-2  
administrative procedures 5-1  
ALTER\_HIGH\_WATER\_MARK 5-14  
altering partitions A-3  
altpart.sql A-3  
ANYINTERACT 7-8  
arcs A-10  
area Glossary-1  
attribute Glossary-1

## B

---

boundary Glossary-1  
bounded data 9-10  
bounded value 9-3  
BUILD\_WINDOW 8-2  
BUILD\_WINDOW\_FIXED 8-4  
bulk loading 2-2

## C

---

Cartesian Glossary-1  
circles A-10  
CLEAN\_WINDOW 8-6  
consistency check 7-10  
CONTAINS 7-8, Glossary-1  
control file 2-2  
coordinate Glossary-2  
coordinate system 7-2, Glossary-2  
COVEREDBY 7-8  
COVERS 7-8, Glossary-2  
cr\_spatial\_index.sql A-1

CREATE\_WINDOW\_LAYER 8-7  
creating layer tables A-2  
crlayer.sql A-2

## D

---

data Glossary-2  
data dictionary B-1  
data model 1-3  
DATE data type 9-6, 9-11  
decompose Glossary-2  
dimensional Glossary-2  
disjoint 7-6, 7-9, Glossary-2  
DROP\_PARTITION\_INFO 5-15  
dropping partitions A-3  
drppart.sql A-3  
dynamic query window 3-5

## E

---

element 1-3  
ENCLOSES 9-4  
encoding dimensions 9-9  
EQUAL 7-9, 9-4, Glossary-2  
error messages C-1  
ESTIMATE\_TILING\_LEVEL 6-2  
extent 6-5, Glossary-2  
EXTENT\_OF 6-5  
extracting a dimension 9-8

## F

---

feature Glossary-2  
filter 3-7

fixed-size tiles 2-9, 5-4, 5-10

## G

---

Geographical Information System Glossary-3  
geometric objects 1-3  
geometric primitive 1-2  
georeferenced Glossary-3  
GIS 1-1, Glossary-3  
grants 5-18, A-3  
grid Glossary-3

## H

---

HHCODE Glossary-3  
high water mark 5-14, Glossary-3  
hole in a polygon 2-7, Glossary-3  
homogeneous Glossary-3  
hyperspatial Glossary-4

## I

---

index 5-2, 5-4, 5-8, 5-10, Glossary-4  
index creation 2-8  
INIT\_ELEMENT 7-4  
inserting spatial data 2-4  
INSIDE 7-9, 9-4  
INTERACT 7-5  
interaction 7-8

## K

---

key Glossary-4  
keyword Glossary-4

## L

---

latitude Glossary-4  
layer 1-4, A-2  
line Glossary-4  
line data 1-3  
loading process 2-2  
location 1-1  
longitude Glossary-4

## M

---

minimum bounding rectangle 6-5

## O

---

OUTSIDE 9-4  
OVERLAP 9-4  
OVERLAPBDYDISJOINT 7-9  
OVERLAPBDYINTERSECT 7-9

## P

---

partition 5-16, Glossary-4, Glossary-5  
partition key 4-1  
partitioned table 4-1, 5-15, 5-16, 5-19, A-3,  
Glossary-5  
partitioned tables 1-17  
PL/SQL 9-1  
plotting tiles A-11  
point data 1-3, 4-1  
polygon Glossary-5  
polygon data 1-3  
POPULATE\_INDEX 5-2  
POPULATE\_INDEX\_FIXED 5-4  
primary filter 3-7  
primitive 1-2  
PROPAGATE\_GRANTS 5-18  
proximity Glossary-5

## Q

---

query Glossary-5  
query window 3-5, Glossary-5

## R

---

RDBMS Glossary-5, Glossary-6  
recursion Glossary-5  
region Glossary-5  
REGISTER\_PARTITION\_INFO 5-19  
RELATE 7-7  
REPARTITION 5-20  
resolution Glossary-6

## S

---

scale Glossary-6  
schema 1-4  
SD\*Converter Glossary-6  
SD\*Loader Glossary-6  
SDO\_BVALUETODIM 9-3  
SDO\_CODE\_SIZE 5-7  
SDO\_COMPARE 9-4  
SDO\_DATETODIM 9-6  
SDO\_DECODE 9-8  
SDO\_ENCODE 9-9  
SDO\_TO\_BVALUE 9-10  
SDO\_TO\_DATE 9-11  
sdogrant.sql A-3  
secondary filter 3-8  
server partitioning 4-1  
SLF Glossary-6  
sort Glossary-6  
Spatial Glossary-7  
spatial Glossary-6, Glossary-7  
spatial data model Glossary-7  
spatial database Glossary-6  
spatial index 1-9, 2-8, 5-8, 5-10  
spatial join 3-9  
Spatial Load Format Glossary-7  
spatial query 3-5, Glossary-7  
spatiotemporal Glossary-7  
SQL script A-1, A-2  
SQL\*Loader 2-2

## T

---

table partitioning 1-17, 4-1  
tessellation 1-9, 2-8, 5-2, 5-4, 5-8, 5-10,  
Glossary-7  
tile 1-9, 3-2  
tiling 5-10, 6-2, A-4, Glossary-7  
TOUCH 7-9, Glossary-7  
transactional insert 2-4  
two-tier query 3-1

## U

---

UPDATE\_INDEX 5-8  
UPDATE\_INDEX\_FIXED 5-10

## V

---

VALIDATE\_GEOMETRY 7-10  
VERIFY\_LAYER 5-12  
VERIFY\_PARTITIONS 5-22  
visualizing tiles A-11