



Chapter 24: Advanced Data Types and New Applications

- Temporal Data
- Spatial and Geographic Databases
- Multimedia Databases
- Mobility and Personal Databases





Time In Databases

- While most databases tend to model reality at a point in time (at the "current" time), temporal databases model the states of the real world across time.
- The **valid time** is the set of time intervals during which the fact is true in the real world.
- The **transaction time** for a fact is the time interval during which the fact is current within the database system.
- In a **temporal relation**, each tuple has an associated time when it is true; the time may be either valid time or transaction time.
- A **bi-temporal relation** stores both valid and transaction time.





Time In Databases (Cont.)

- Example of a temporal relation:

<i>account-number</i>	<i>branch-name</i>	<i>balance</i>	<i>from</i>	<i>to</i>
A-101	Downtown	500	1999/1/1 9:00	1999/1/24 11:30
A-101	Downtown	100	1999/1/24 11:30	*
A-215	Mianus	700	2000/6/2 15:30	2000/8/8 10:00
A-215	Mianus	900	2000/8/8 10:00	2000/9/5 8:00
A-215	Mianus	700	2000/9/5 8:00	*
A-217	Brighton	750	1999/7/5 11:00	2000/5/1 16:00

- “*” means the tuple is true until the *to* field is written: this means the tuple is true at current time.
- Temporal query languages have been proposed to simplify modeling of time as well as time related queries.





Time Specification in SQL-92

- **date**: four digits for the year (1--9999), two digits for the month (1--12), and two digits for the date (1--31).
- **time**: two digits for the hour, two digits for the minute, and two digits for the second, plus optional fractional digits.
- **timestamp**: the fields of **date** and **time**, with six fractional digits for the seconds field.
- Times are specified in the *Universal Coordinated Time*, abbreviated UTC (from the French); supports **time with time zone**.
- **interval**: refers to a period of time (e.g., 2 days and 5 hours), without specifying a particular time when this period starts; could more accurately be termed a *span*.





Temporal Query Languages

- Predicates *precedes*, *overlaps*, and *contains* on time intervals.
- *Intersect* can be applied on two intervals, to give a single (possibly empty) interval; the union of two intervals may or may not be a single interval.
- A **snapshot** of a temporal relation at time t consists of the tuples that are valid at time t , with the time-interval attributes projected out.
- **Temporal selection**: involves time attributes
- **Temporal projection**: the tuples in the projection **inherit** their time-intervals from the tuples in the original relation.
- **Temporal join**: the time-interval of a tuple in the result is the **intersection** of the time-intervals of the tuples from which it is derived. If intersection is empty, tuple is discarded from join.





Spatial and Geographic Databases

Database System Concepts

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use





Spatial and Geographic Databases

- Spatial databases store information related to spatial locations, and support efficient storage, indexing and querying of spatial data.
- Special purpose index structures are important for accessing spatial data, and for processing spatial join queries.
- **Computer Aided Design (CAD)** databases store design information about how objects are constructed E.g.: designs of buildings, aircraft, layouts of integrated-circuits
- Geographic databases store geographic information (e.g., maps): often called **geographic information systems or GIS**.
- Products
 - IBM DB2 Spatial Extender
 - Informix Spatial Datablade
 - Oracle Spatial





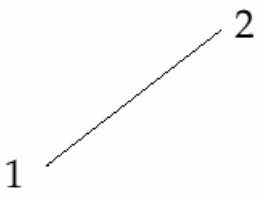
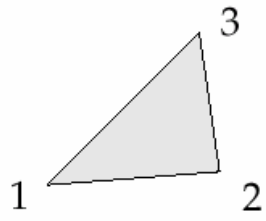
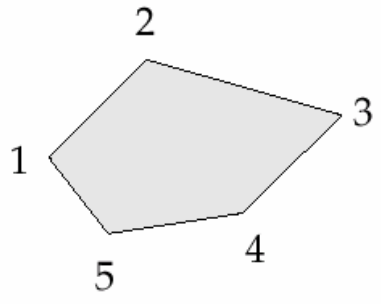
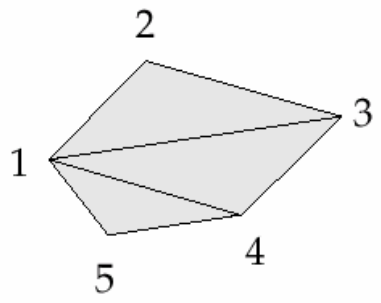
Representing Geometric Information

- Various geometric constructs can be represented in a database in a normalized fashion.
- Represent a line segment by the coordinates of its endpoints.
- Approximate a curve by partitioning it into a sequence of segments
 - Create a list of vertices in order, or
 - Represent each segment as a separate tuple that also carries with it the identifier of the curve (2D features such as roads).
- Closed polygons
 - List of vertices in order, starting vertex is the same as the ending vertex, or
 - Represent boundary edges as separate tuples, with each containing identifier of the polygon, or
 - Use **triangulation** — divide polygon into triangles
 - ▶ Note the polygon identifier with each of its triangles.





Representation of Geometric Constructs

line segment		$\{(x_1, y_1), (x_2, y_2)\}$
triangle		$\{(x_1, y_1), (x_2, y_2), (x_3, y_3)\}$
polygon		$\{(x_1, y_1), (x_2, y_2), (x_3, y_3), (x_4, y_4), (x_5, y_5)\}$
polygon		$\{(x_1, y_1), (x_2, y_2), (x_3, y_3), ID1\}$ $\{(x_1, y_1), (x_3, y_3), (x_4, y_4), ID1\}$ $\{(x_1, y_1), (x_4, y_4), (x_5, y_5), ID1\}$
	object	representation





Representation of Geometric Information (Cont.)


- Representation of points and line segment in 3-D similar to 2-D, except that points have an extra z component
- Represent arbitrary polyhedra by dividing them into tetrahedrons, like triangulating polygons.
- Alternative: List their faces, each of which is a polygon, along with an indication of which side of the face is inside the polyhedron.






Design Databases

- CAD applications
 - Aeroplane design: impossible to keep the complete design in memory
 - Solution: OODB
- OODB represent design components as objects (generally geometric objects): the connections between the objects indicate how the design is structured.
- Simple two-dimensional objects: points, lines, triangles, rectangles, polygons.
- Complex two-dimensional objects: formed from simple objects via union, intersection, and difference operations.
- Complex three-dimensional objects: formed from simpler objects such as spheres, cylinders, and cuboids, by union, intersection, and difference operations.
- Wireframe models represent three-dimensional surfaces as a set of simpler objects.






Representation of Geometric Constructs



(a) Difference of cylinders

(b) Union of cylinders

- Design databases also store non-spatial information about objects (e.g., construction material, color, etc.)
- **Spatial integrity constraints** are important.
 - E.g., pipes should not intersect, wires should not be too close to each other, etc.





Geographic Data

- **Raster data** consist of bit maps or pixel maps, in two or more dimensions.
 - Example 2-D raster image: satellite image of cloud cover, where each pixel stores the cloud visibility in a particular area.
 - Additional dimensions might include the temperature at different altitudes at different regions, or measurements taken at different points in time.
- Design databases generally do not store raster data.





Geographic Data (Cont.)

- **Vector data** are constructed from basic geometric objects: points, line segments, triangles, and other polygons in two dimensions, and cylinders, spheres, cuboids, and other polyhedrons in three dimensions.
- Vector format often used to represent map data.
 - Roads can be considered as two-dimensional and represented by lines and curves.
 - Some features, such as rivers, may be represented either as complex curves or as complex polygons, depending on whether their width is relevant.
 - Features such as regions and lakes can be depicted as polygons.





Applications of Geographic Data

- Examples of geographic data
 - map data for vehicle navigation
 - distribution network information for power, telephones, water supply, and sewage
- Vehicle navigation systems store information about roads and services for the use of drivers:
 - **Spatial data:** e.g, road/restaurant/gas-station coordinates
 - **Non-spatial data:** e.g., one-way streets, speed limits, traffic congestion
- **Global Positioning System (GPS)** unit - utilizes information broadcast from GPS satellites to find the current location of user with an accuracy of tens of meters.
 - increasingly used in vehicle navigation systems as well as utility maintenance applications.





Spatial Queries

- **Nearness queries** request objects that lie near a specified location.
- **Nearest neighbor queries**, given a point or an object, find the nearest object that satisfies given conditions.
 - Locate faulty elements in complex design projects
- **Region queries** deal with spatial regions. e.g., ask for objects that lie partially or fully inside a specified region.
- Queries that compute intersections or **unions** of regions.
- **Spatial join** of two spatial relations with the location playing the role of join attribute.





Spatial Queries (Cont.)

- Spatial data is typically queried using a **graphical query language**; results are also displayed in a graphical manner.
- Graphical interface constitutes the front-end
- Extensions of SQL with abstract data types, such as lines, polygons and bit maps, have been proposed to interface with back-end.
 - allows relational databases to store and retrieve spatial information
 - Queries can use spatial conditions (e.g. contains or overlaps).
 - queries can mix spatial and nonspatial conditions





Multimedia Databases

Database System Concepts

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use





Multimedia Databases

- One approach: store in files
- Another approach: store attributes in databases and the multimedia objects in files
- To provide database functions such as indexing and consistency, it is desirable to store multimedia data in a database
- The database must handle large object representation.
 - Some DBs cannot handle objects of more than a few GBs
- Similarity-based retrieval must be provided by special index structures.
 - Fingerprint database
- Must provide guaranteed steady retrieval rates for continuous-media data.
 - Audio or video data





Multimedia Data Formats

- Store and transmit multimedia data in compressed form
 - JPEG and GIF the most widely used formats for image data.
 - MPEG standard for video data use commonalities among a sequence of frames to achieve a greater degree of compression.
- MPEG-1 quality comparable to VHS video tape.
 - stores a minute of 30-frame-per-second video and audio in approximately 12.5 MB
- MPEG-2 designed for digital broadcast systems and digital video disks; negligible loss of video quality.
 - Compresses 1 minute of audio-video to approximately 17 MB.
- Several alternatives of audio encoding
 - MPEG-1 Layer 3 (MP3), RealAudio, WindowsMedia format, etc.





Continuous-Media Data

- Most important types are video and audio data.
- Characterized by high data volumes and real-time information-delivery requirements.
 - Data must be delivered **sufficiently fast** that there are no gaps in the audio or video.
 - Data must be delivered at a rate that does not cause **overflow** of system buffers.
 - **Synchronization** among distinct data streams must be maintained
 - ▶ video of a person speaking must show lips moving synchronously with the audio





Video Servers

- **Video-on-demand** systems deliver video from central video servers, across a network, to terminals
 - Must guarantee end-to-end delivery rates
- Current video-on-demand servers are based on **file systems**; existing database systems do not meet real-time response requirements.
- Multimedia data are stored on several disks (RAID configuration), or on tertiary storage for less frequently accessed data.
- Head-end terminals - used to view multimedia data
 - PCs or TVs attached to a small, inexpensive computer called a set-top box.






Similarity-Based Retrieval

Examples of similarity based retrieval

- Pictorial data: Two pictures or images that are slightly different as represented in the database may be considered the same by a user.
 - E.g., identify similar designs for registering a new trademark.
- Audio data: Speech-based user interfaces allow the user to give a command or identify a data item by speaking.
 - E.g., test user input against stored commands.
- Handwritten data: Identify a handwritten data item or command stored in the database





Mobility

Database System Concepts

©Silberschatz, Korth and Sudarshan


See www.db-book.com for conditions on re-use





Mobile Computing Environments

- A mobile computing environment consists of mobile computers, referred to as **mobile hosts**, and a wired network of computers.
- Mobile host may be able to communicate with wired network through a *wireless digital communication network*
 - Wireless local-area networks (within a building)
 - ▶ E.g. Avaya's Orinico Wireless LAN
 - Wide areas networks
 - ▶ Cellular digital packet networks
 - 3 G and 2.5 G cellular networks





Mobile Computing Environments (Cont.)

- A model for mobile communication
 - Mobile hosts communicate to the wired network via computers referred to as **mobile support (or base) stations**.
 - Each mobile support station manages those mobile hosts within its **cell (geographical area it covers)**
 - When mobile hosts move between cells, there is a **handoff** of control from one mobile support station to another.
- Direct communication, without going through a mobile support station is also possible between nearby mobile hosts
 - Supported, for e.g., by the Bluetooth standard (up to 10 meters, at up to 721 kbps)





Database Issues in Mobile Computing


- New issues for query optimization.
 - Connection time monetary charges and number of bytes transmitted
 - Energy (battery power) is a scarce resource and its usage must be minimized
- Mobile user's locations may be a parameter of the query
 - GIS queries
 - Techniques to track locations of large numbers of mobile hosts
- Broadcast data can enable any number of clients to receive the same data at no extra cost
 - leads to interesting querying and data caching issues.
- Users may need to be able to perform database updates even while the mobile computer is disconnected.
 - e.g., mobile salesman records sale of products on (local copy of) database.
 - Can result in conflicts detected on reconnection, which may need to be resolved manually.





Routing and Query Processing


- Must consider these competing costs:
 - User time.
 - Communication cost
 - ▶ Connection time - used to assign monetary charges in some cellular systems.
 - ▶ Number of bytes, or packets, transferred - used to compute charges in digital cellular systems
 - ▶ Time-of-day based charges - vary based on peak or off-peak periods
 - Energy - optimize use of battery power by minimizing reception and transmission of data.
 - ▶ Receiving radio signals requires much less energy than transmitting radio signals.





Broadcast Data

- Mobile support stations can broadcast frequently-requested data
 - Allows mobile hosts to wait for needed data, rather than having to consume energy transmitting a request
 - Supports mobile hosts without transmission capability
- A mobile host may optimize energy costs by determining if a query can be answered **using only cached data**
 - If not then must either;
 - ▶ Wait for the data to be broadcast
 - ▶ Transmit a request for data and must know when the relevant data will be broadcast.
- Broadcast data may be transmitted according to a fixed schedule or a changeable schedule.
 - For changeable schedule: the broadcast schedule must itself be broadcast at a well-known radio frequency and at well-known time intervals
- Data reception may be interrupted by noise
 - Use techniques similar to RAID to transmit redundant data (parity)





End of Chapter

Database System Concepts

©Silberschatz, Korth and Sudarshan

See www.db-book.com for conditions on re-use

