CLAIMS

1. A database system configured to search plural mobile objects, each object having attributes including location data, to determine neighbouring objects nearest to a specific location, said object being located in a geographical space made up of plural spatially distinct subspaces, each being made up of plural cells, the database system comprising

plural storage nodes; and

an operating system, the operating system being configured to control storage of object data amongst the plural storage nodes,

wherein the operating system is configured to cause storage of data representative of one or more spatially distinct subspaces in a respective single one of the plural storage nodes,

wherein location data of each object is used to index that object with respect to cells making up each spatially distinct subspace in each node, and

wherein data are stored in the plural storage nodes using a configurable mapping from subspaces to storage nodes which explicitly defines which subspace belongs to which storage node, based on a read and/or write load on each of the spatially distinct subspaces.

2. A database system according to claim 1, wherein the data of each spatially distinct subspace is stored completely in a single storage node.

3. A database system according to claim 1 or 2, wherein the operating system is configured such that data of each spatially distinct subspace is replicated to plural storage nodes to form data replicas.

4. A database system according to claim 3 wherein the operating system is configured such that write operations concerning a spatially distinct subspace are propagated to all the relevant data replicas.

5. A database system according to claim 3 or 4 wherein the number of replicas is configurable based on use cases.

6. A database system according to any preceding claim wherein the operating system is configured to operate a breadth-first search algorithm to answer K-nearest neighbour queries

7. A database system according to claim 1, wherein data are stored in the plural storage nodes by consistent hashing.

8. A database system according to claim 1, wherein for load balancing, the operating system is configured to use both a user-configurable mapping from subspaces to storage nodes which explicitly defines which subspace belongs to which node, and consistent hashing.

9. A database system according to claim 8, wherein for data not included in the mapping, consistent hashing is employed.

10. A database system according to claim 1, in which one node in the mapping is used as a static coordinator to broadcast new joins.

11. A database system according to claim 1, wherein the operating system applies gossip style messaging for node discovery.

12. A database system for a ride hailing application according to any preceding claim wherein the objects are service provider vehicles.

13. A method of storing data representing plural mobile objects, each object having attributes including location data, to enable fast searching for neighbours nearest to a specific location in a geographical space made up of plural spatially distinct subspaces, each being made up of plural cells, the database system comprising plural storage nodes; the method comprising:

storing object data amongst the plural storage nodes, such that data representative of one or more spatially distinct subspaces is stored in a respective single one of the storage nodes,

using current location data of each object to index that object with respect to cells making up each spatially distinct subspace in each storage node, and

using a read and/or write load on each of the spatially distinct subspaces to map the subspaces to storage nodes to explicitly define which subspace belongs to which storage node.

14. A scalable in-memory spatial data store for kNN searches comprising a database system as claimed in any of claims 1-12.