

CAN

Zheng Ma

Fall 2003

Content-Addressable-Network (CAN)

▪ CAN: Internet Scale Hash table

▪ Interface

▪ insert(key,value)

▪ value = retrieve(key)

▪ Idea: associate to each node and item a unique coordinate in an d-dimensional Cartesian space.

▪ Properties

-scalable

-operationally simple

-good performance

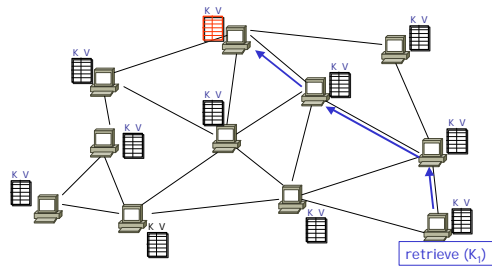
CAN: basic idea

CAN: basic idea (cont.)

CAN: basic idea (cont.)

CAN: basic idea (cont.)

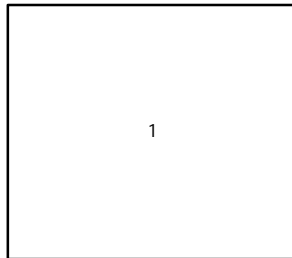
CAN: basic idea (cont.)



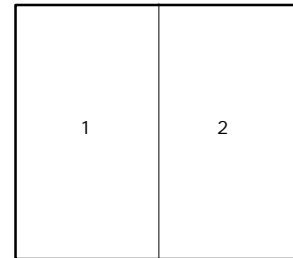
CAN (Solution)

- virtual Cartesian coordinate space
 - entire space is partitioned amongst all the nodes
 - every node "owns" a zone in the overall space
- abstraction
 - can store data at "points" in the space
 - can route from one "point" to another
- point = node that owns the enclosing zone

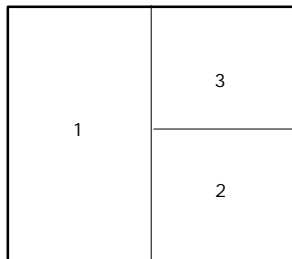
CAN (Simple example 2-dimension)



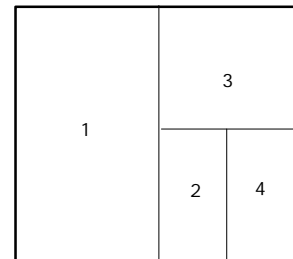
CAN (Simple example cont.)

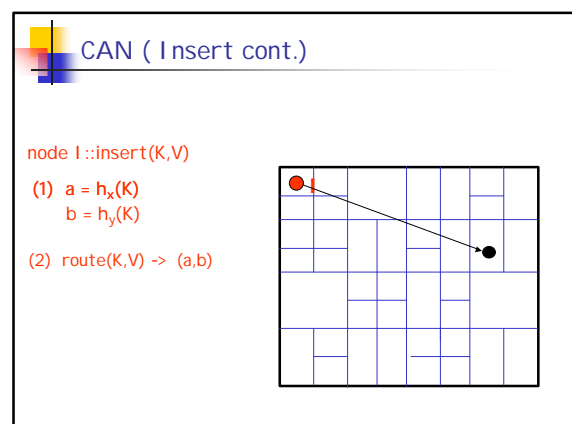
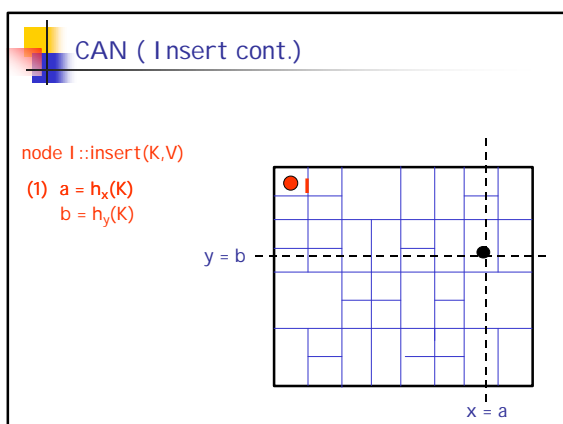
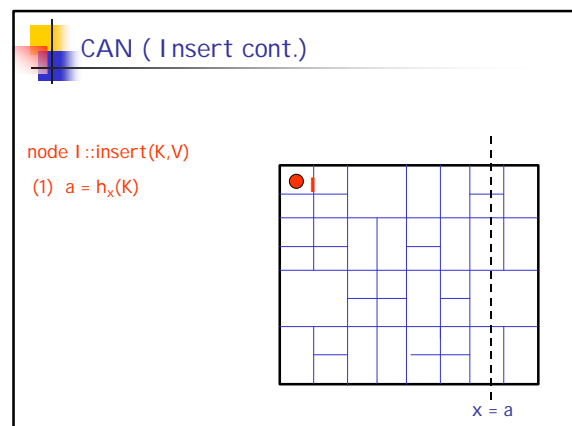
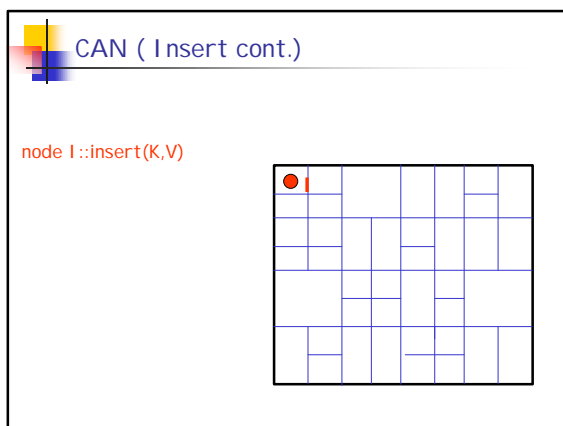
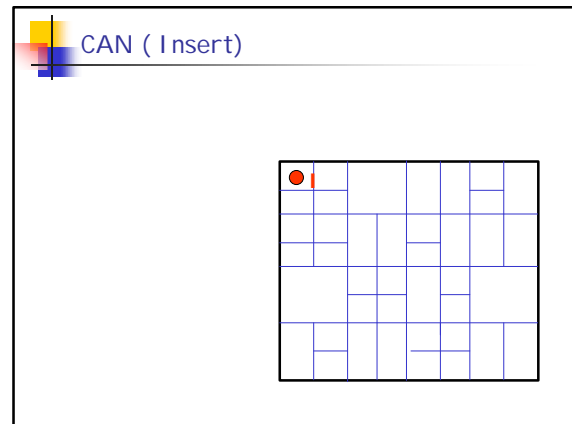
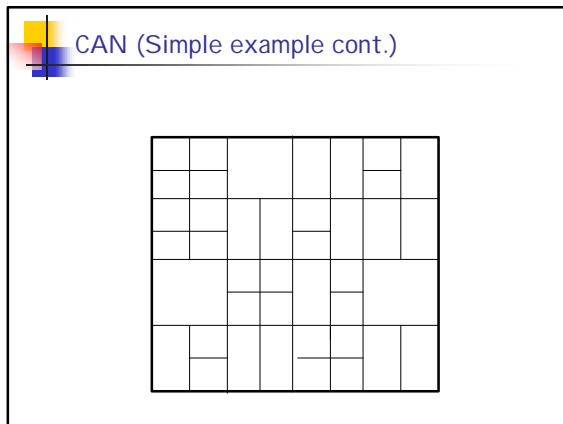


CAN (Simple example cont.)



CAN (Simple example cont.)





CAN (Insert cont.)

node I :: insert(K,V)

- (1) $a = h_x(K)$
 $b = h_y(K)$
- (2) route(K,V) \rightarrow (a,b)
- (3) (a,b) stores (K,V)

CAN (Retrieve)

node J :: retrieve(K)

- (1) $a = h_x(K)$
 $b = h_y(K)$
- (2) route "retrieve(K)" to (a,b)

CAN (A missing part)

Data stored in the CAN is addressed by name (i.e. key), not location (i.e. IP address)

Question: What is missing in the procedure ?

Routing !

CAN (Routing table)

A node only maintains state for its immediate neighboring nodes

CAN (Routing)

CAN (Node Join)

Bootstrap node

new node

1) Discover some node "I" already in CAN

(p,q) \rightarrow 2) pick random point in space

