GHT: A Geographic Hash Table for Data-Centric Storage

Sylvia Ratnasamy, Scott Shenker,
Brad Karp, Ramesh Govindan, Deborah Estrin
ICSI/UCB/USC/UCLA

Based on
Sylvia Ratnasamy's slides

Keywords and Terminology

- Sensor network
  - a distributed network comprised of a large number of small sensing devices equipped with computation, communication and storage
- Observation
  - low-level readings from sensors
  - e.g. Detailed temperature readings
- Events
  - Predefined constellations of low-level observations
  - e.g. temperature greater than 75 F

Total Usage /Hotspot Usage

- Total Usage
  Total number of packets sent in the sensor network
- Hotspot Usage
  The maximal number of packets sent by a particular sensor node

Design Goals

- Availability
- Consistency
- Scaling
- Topological generality

Data-Centric Storage

- Data items are stored by "name" at nodes within the sensor network
- All data with the same general name will be stored at the same node
  e.g. ("elephant sightings")
- Queries for data with a particular name are sent to the node storing those named data

Geographic Hash Table

- Events are named with keys and both the storage and the retrieval are performed using keys
- Hash function maps the key into geographic coordinates

GHT Algorithm

Based on GPSR where packets are marked with position of destinations and each node is aware of its position

- Greedy forwarding algorithm
- Perimeter forwarding algorithm

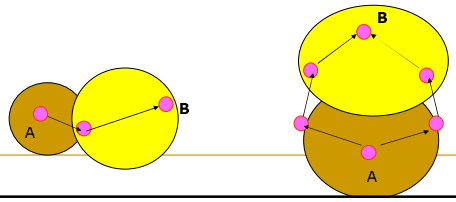

Perimeter Forwarding Algorithm

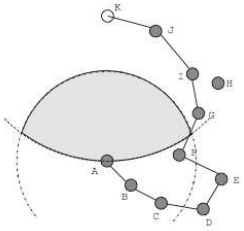

Home Node

- Home node: node closest to destination
- Packet not addressed to specific node but only to a specific location
- Packet traverses entire perimeter that encloses the destination before being consumed by the home node.

Problems

- Robustness could be affected
  - Nodes could move
  - Node failure can occur
  - Deployment of new nodes
- Not Scalable
  - Bottleneck at home nodes

Solutions to the problems

- Perimeter refresh protocol
- Structured replication

Perimeter Refresh Protocol

- Replicates stored data for a key at nodes around the location to which this key hashes.
- Home node will generate refresh packet periodically.
- Refresh packet is addressed to the hashed location of the key.
- If replica doesn't receive refresh packet after some time, it will initiate a refresh.

Structured Replication

- When too many events are detected, home node becomes the hotspot of communication.
- Hierarchical decomposition of the key space.
- Reduced the cost of storage, but increases the cost of queries.

Simulations

- To compare the storage methods in terms of total and hotspot usage

Simulation Results under varying *n*

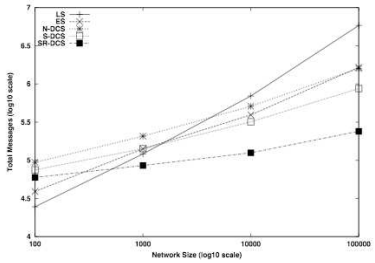

Number of Queries Q = 50

Simulation Results under varying *n*

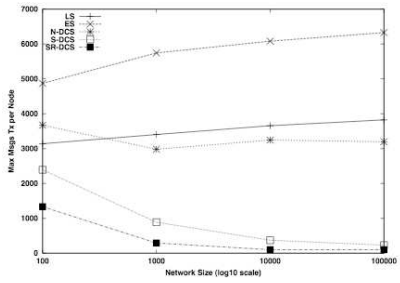

Number of Queries Q = 50

Conclusion

- DCS entails storing data items by name at nodes within the sensor network
- GHT hashes the key (events) into geographic co-ordinates and stores a key-value pair at the sensor node geographically nearest to the hash
- GHT uses Perimeter Refresh Protocol and Structured Replication to achieve robustness and scalability
- DCS appears to scale well to large sensor networks

REFERENCES

- Deepak Ganesan, Deborah Estrin, John Heidemann, Dimensions: why do we need a new data handling architecture for sensor networks?, ACM SIGCOMM Computer Communication Review, Volume 33 Issue 1, January 2003 Scott Shenker, Sylvia Ratnasamy, Brad Karp, Ramesh Govindan, Deborah Estrin, Data-centric storage in sensornets, ACM SIGCOMM Computer Communication Review, Volume 33 Issue 1, January 2003
- Sylvia Ratnasamy, Brad Karp, Scott Shenker, Deborah Estrin, Ramesh Govindan, Li Yin, Fang Yu, Data-centric storage in sensornets with GHT, a geographic hash table, Mobile Networks and Applications, Volume 8 Issue 4, August 2003
- Chalermek Intanagonwiwat, Ramesh Govindan, Deborah Estrin, John Heidemann, Fabio Silva, Directed diffusion for wireless sensor networking, IEEE/ACM Transactions on Networking (TON), Volume 11 Issue, February 2003
- R. Govindan, J. M. Hellerstein, W. Hong, S. Madden, M. Franklin, S. Shenker, The Sensor Network as a Database, USC Technical Report No. 02-771, September 2002