Time Synchronization and Location Discovery in Sensor Networks

Leonidas Guibas
Stanford University

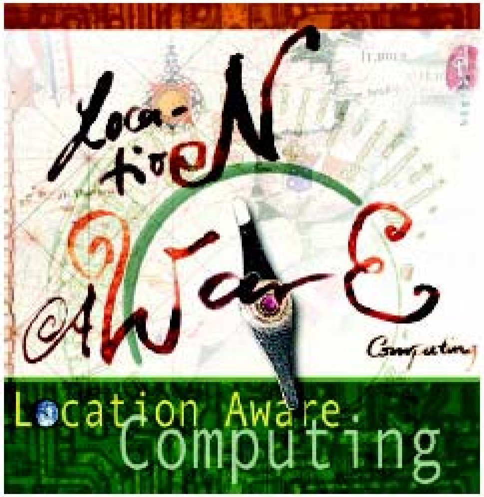

Infrastructure Establishment

- Sensor networks are typically deployed in an *ad hoc* setting
- Before they can usefully function, sensor nodes must collaborate to establish a common temporal and spatial frame of reference:
  - *time synchronization*
    - to a global clock
    - with mappings between node clocks
  - *localization in physical space*
    - using beacon nodes (say with GPS)
    - exploiting multiple distance measurements to other nodes
  - *location services*

Papers for Today

Synchronization

[roemer01] Kay Römer. "Time Synchronization in Ad Hoc Networks." Proc. 2001 ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc) 2001, pp. 173-182.

[elson02] J. Elson, L. Girod, and D. Estrin. "Fine Grained Network Time Synchronization using Reference Broadcasts." Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI 2002), Boston, Massachusetts, December, 2002.

Localization (Discovery, Services)

[li00] Jinyang Li, John Jannotti, Douglas S. J. De Couto, David R. Karger, and Robert Morris. "A Scalable Location Service for Geographic Ad Hoc Routing." ACM Mobicom 2000, Boston, MA, pages 120-130.

[savvides03] Andreas Savvides, and Mani B. Strivastava. "Distributed Fine-Grained Localization in Ad-Hoc Networks." IEEE Transactions of Mobile Computing, to appear, 2003.

Time Synchronization

- Essential for temporal reasoning (is the car moving north or south?)
- Essential for distance or velocity estimation of moving objects
- Essential for many algorithms (node localization, target tracking, etc.)

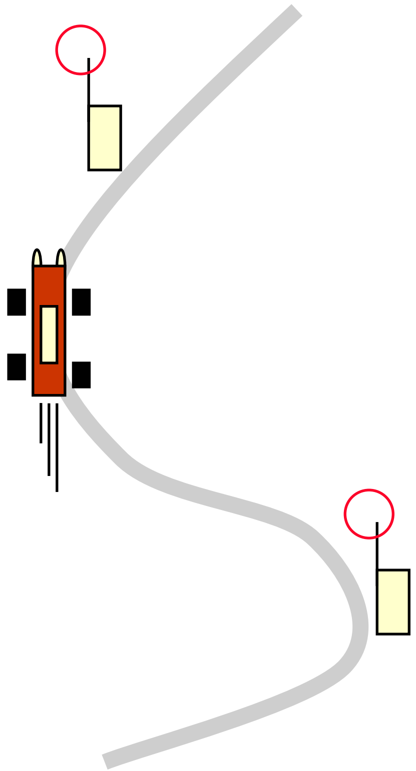

Need to perform time comparisons and estimate time differences

Synchronization Challenges in Sensor Networks

- Links can go up and down; the network may often be partitioned
- Communication delay hard to estimate, even for a single hop. The delay components are:
  - send time
  - access time
  - propagation time
  - receive time
- Highly variable communication delay between pairs of nodes requiring multiple hops
- Clock skew among nodes is common
- Usually no special node with accurate clock (UTC)

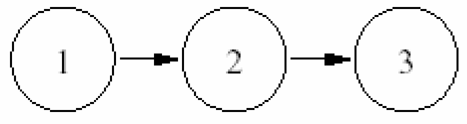

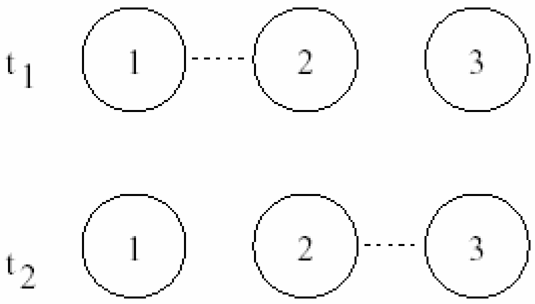

[Römer, 2001]: Interval Arithmetic on Local Times

- Generate time stamps using unsynchronized local clocks with known bounds on skew
- Use interval methods to estimate and bound message delays
  - Exploit ACK messages, inter-message durations
- Translate the time stamp to an interval in the local time of the receiving node
- Use intervals for temporal reasoning

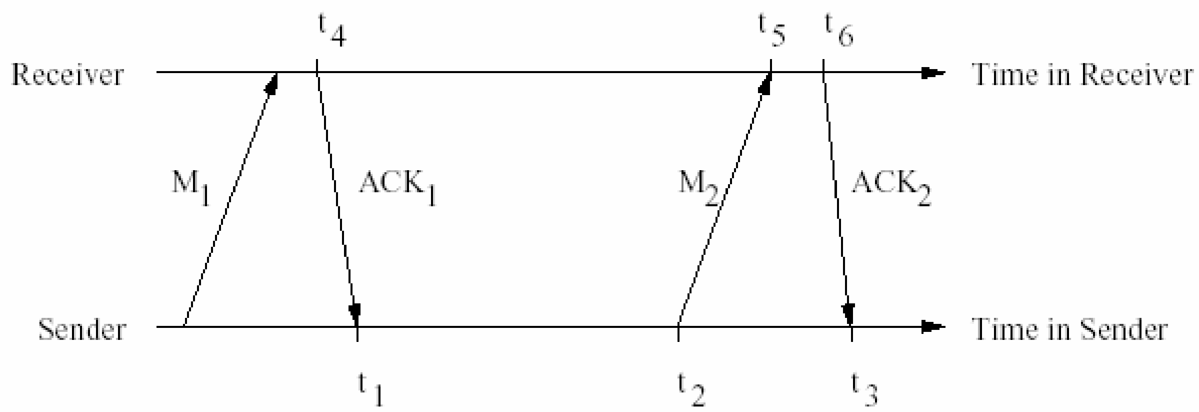

[Elson, Girod, and Estrin, 2002]: Reference Broadcasts for Receiver Synchronization

- Exploit broadcast nature of radio communication
- Use a sender *to synchronize multiple receivers* (inside the sender's communication range) *to each other*
  - thus removing *send* and *access time* variability
  - furthermore, propagation time is usually negligible, and receive time can be more tightly controlled
- Use repeated reference broadcasts between two receivers to eliminate outlier data and do a least-squares fit that maps from one receiver clock to the other (estimating both relative clock phase offset and skew)

Location Discovery

Knowing node locations is important for:

- routing and communication (GPSR, GeoRouting)
- sensor network organization (sensor collaboration groups, sentry selection)
- information about locally available resources (specialized stationary nodes, recharging stations, etc.)
- more meaningful data collection (data tagged with location, redundancy suppression)
- adaptation to the environment (stealth, location-aware applications)

Taxonomy of Location Discovery Systems

- Absolute/relative
- Localized/centralized
- Accuracy/precision
- Scale
- Cost

- GPS
- Active Badges (Olivetti)
- Active Bats (AT&T)
- Cricket (MIT)
- RADAR (Microsoft)
- EasyLiving (Microsoft)
- SmartFloor (Georgia Tech)
- AHLoS (UCLA)

Questions to Keep in Mind

- What is the physical basis of the location discovery method?
- Are the algorithms more of the client-server, or the peer-to-peer type?
- How does the system dear with interference, incomplete data, etc?
- Is there a fixed infrastructure required (say beacons)? What is its cost? How easily is it deployed?

[Savvides and Srivastava, 2003]: Node Localization by Distributed Non-Linear Optimization

- Sensor network is a mixture of beacons and nodes of unknown position
- Radio signals or ultrasonic sensors can be used to estimate distance between nodes
- Need to determine the positions of the non-beacon nodes, using this information.

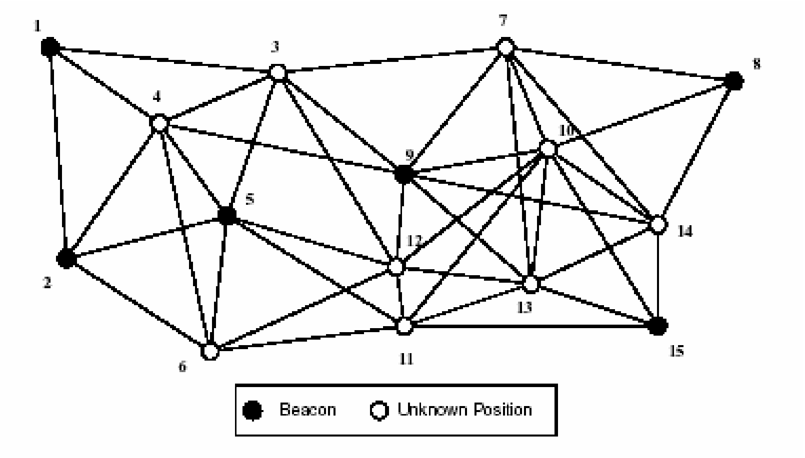

Atomic and Iterative Multilateration

- Determination of the position of a node, using distance measurements to beacons – at least three are required
- If that is not possible, then use iterative refinement
- Distributed implementation is possible

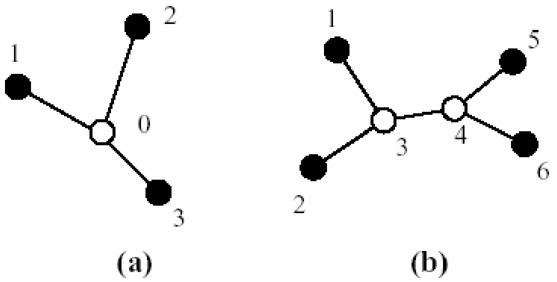

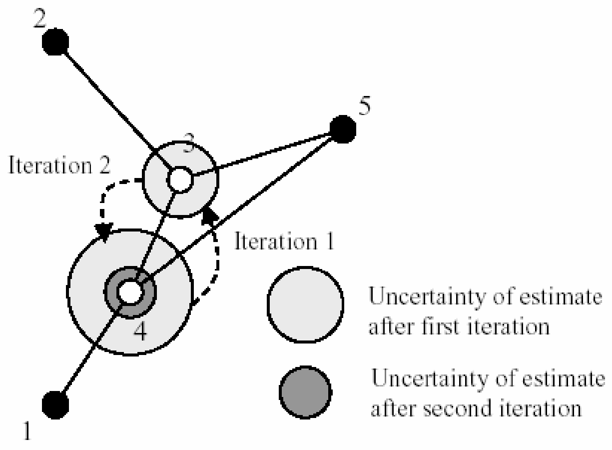

Location Services

- Geographic routing requires a node to know the location of a message destination
- After location discovery, we have to decide who knows the location of whom
  - full exchange of location information is an expensive $O(n^2)$ process
  - but without some preprocessing to make location information more easily accessible, geographic routing is impossible
- Thus the need for location services; also useful as nodes move, get depleted and die, or are added to the network

[Li *et. al.*, 2000]: Hierarchical Technique for Allocating Location Servers (GRID)

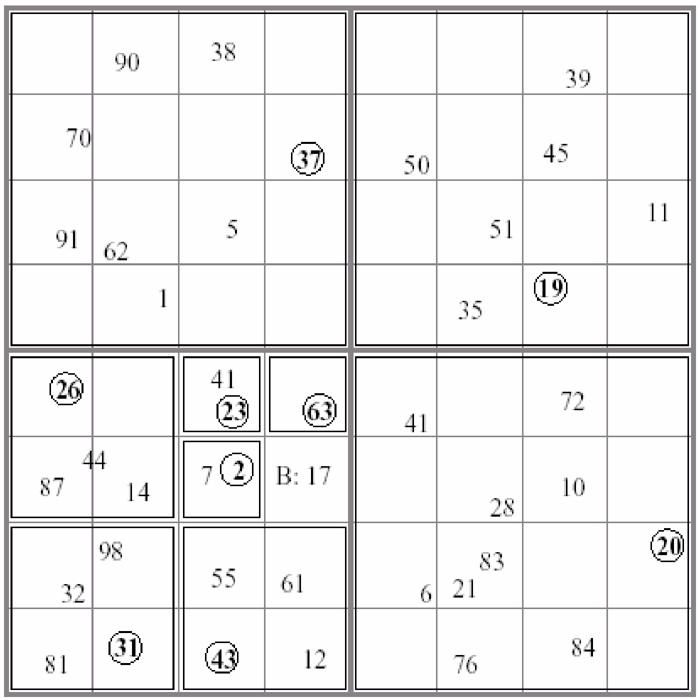

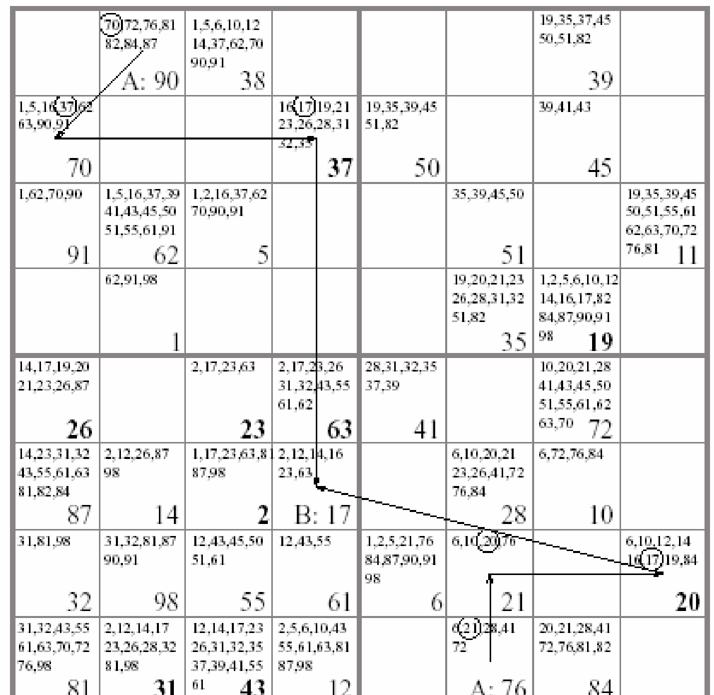

GRID Features

- Each node acts as a location server.
- If node IDs are randomly assigned, a node will typically be a location server for O(log*n*) other nodes.
- If the lowest common ancestor of two nodes *u* and *v* is at level *k* of the quad-tree hierarchy, GRID will answer a query from *u* for *v*'s location in O(*k*) steps [distance sensitive algorithm].

Conclusion

- Synchronization and localization are critical processes for enabling functionality in a sensor network.
- Only preliminary work is available in this area – a lot remains to be done.