

18-759: Wireless Networks

Lecture 13: Sensor Networks

Peter Steenkiste
Departments of Computer Science and
Electrical and Computer Engineering
Spring Semester 2010
<http://www.cs.cmu.edu/~prs/wirelessS10/>

Peter A. Steenkiste, CMU

1

Outline

- What technology for what applications?
- WSNs characteristics and design issues, with special focus on:
 - » Power management
 - » Reliable data collection
 - » Hybrid architectures
- Are there size limitations?
- Conclusion

- Based on slides by Prof JP Hubaux, EPFL, and Dr. Lama Nachman, Intel

Peter A. Steenkiste, CMU

2

Cold Chain Management

- Supermarket chains need to track the storage temperature of perishable goods in their warehouses and stores.
- Tens if not hundreds of fridges should be monitored in real-time
- Whenever the temperature of a monitored item goes above a threshold
 - » An alarm is raised and an attendant is warned (pager, sms)
 - » The refrigeration system is turned on
- History of data is kept in the system for legal purpose
- Similar concept can be applied to pressure and temperature monitoring in
 - » Production chains, containers, pipelines




Peter A. Steenkiste, CMU

3

Home automation

- Temperature management
 - » Monitor heating and cooling of a building in an integrated way
 - » Temperature in different rooms is monitored centrally
 - » A power consumption profile is to be drawn in order to save energy in the future
- Lighting management:
 - » Detect human presence in a room to automatically switch lights on and off
 - » Responds to manual activation/deactivation of switches
 - » Tracks movement to anticipate the activation of light-switches on the path of a person
- Similar concept can be applied to
 - » Intrusion detection



Peter A. Steenkiste, CMU

4

Precision Agriculture management

- Farming decisions depend on environmental data (typically photo-synthesis):
 - » Solar radiation
 - » Temperature
 - » Humidity
 - » Soil moisture
- These data evolve continuously over time and space
- A farmer's means of action to influence crop yield :
 - » Irrigation
 - » Fertilization
 - » Pest treatment
- To be optimal, these actions should be highly localized (homogenous parcels can be as small as one hectare or less)
- Environmental impact is also to be taken into account
 - » Salinization of soils, groundwater depletion, well contamination, etc.



Peter A. Steenkiste, CMU

5

WaterSense

- Goal: Help define and implement farming strategies for farmers in a situation of water scarcity.
 - » Crop assessment
 - » Water conservation measures
 - » Time of farming operations
 - » Real-time monitoring of the field conditions
- Desired Outcome: farming decision support system based on environmental data



Peter A. Steenkiste, CMU

6

Earthquake detection

- The occurrence of an earthquake can be detected automatically by accelerometers
- Earthquake speed: around 5-10km/s
- If the epicenter of an earthquake is in an unpopulated area 200km from a city center, instantaneous detection can give a warning up to 30 sec before the shockwave hits the city
- If a proper municipal actuation network is in place:
 - » Sirens go off
 - » Traffic lights go to red
 - » Elevators open at the nearest floor
 - » Pipeline valves are shut
- Even a warning of a few seconds, can reduce the effects of the earthquake
- Similar concept can be applied to
 - » Forest fire, landslides, etc.




Peter A. Steenkiste, CMU

7

Ubiquitous Computing Economic Forecast

- Industrial Monitoring (35% – 45%)
 - Monitor and control production chain
 - Storage management
 - Monitor and control distribution
- Building Monitoring and Control (20 – 30%)
 - Alarms (fire, intrusion etc.)
 - Access control
- Home Automation (15 – 25%)
 - Energy management (light, heating, AC etc.)
 - Remote control of appliances
- Automated Meter Reading (10-20%)
 - Water meter, electricity meter, etc.
- Environmental Monitoring (5%)
 - Agriculture
 - Wildlife monitoring



Other areas:

- Performance monitoring in sports
- Patient monitoring in health and medicine
- Sensor going wireless in vehicular networks

Peter A. Steenkiste, CMU

8

WSN Characteristics and Design Issues

- **Characteristics**
 - » Distributed data collection
 - » Many-to-one (rarely peer-to-peer)
 - » Limited mobility
 - » Data collection (time and space resolution)
 - » Event detection
 - » Minimal intrusiveness
- **Design issues**
 - » Low-cost (hardware and communication)
 - » Extended life-time
 - » Reliable communication
 - » Efficient integrated data processing
 - » Hybrid network infrastructure
 - » Security


Wireless helps
but may not
be required in
all cases!

Peter A. Steenkiste, CMU

9

Wireless Sensor Network architecture


- **Numerous sensor devices**
 - » Modest wireless communication, processing, memory capabilities
 - » Form Ad Hoc Network (self-organized)
 - » Report the measured data to the user



Peter A. Steenkiste, CMU

10

Example of a Low Power Transceiver: Tynynode™




- 868 MHz multi-channel transceiver
- 8 MHz μ -Controller
- 10KB RAM
- 48 kB Program space
- 512 External Flash
- 115 kbps data rate
- 3 V supply voltage
- Current consumption
 - » Transmit 33 mA
 - » Receive 14 mA
 - » Sleep < μ A
- -121 dBm sensitivity
- Radio range 200m (outdoor)
- 39 MHz quartz reference

Peter A. Steenkiste, CMU

11

Sensor Node architecture



- A sensor node can be an *information source*, a *sink* and a *router*
- Autonomous \Rightarrow *low-power*
- Combine *sensing*, *signal conditioning*, *signal processing*, *control* and *communication* capabilities

(courtesy of Swiss Center for Electronics and Microelectronics, Neuchâtel)
Peter A. Steenkiste, CMU

12

Design Issue: Low-cost

- **Hardware**
 - » Low-cost radio
 - » Low cost internal clock
 - » Limited storage and processing capabilities
 - » Not tamper-proof
- **Communication**
 - » Cannot rely on existing pay-per-use cellular infrastructure

Peter A. Steenkiste, CMU

13


Design Issue: Power Management

- **Energy-efficient routing**
 - » Minimum-cost spanning tree
- **Load-balancing**
 - » Mobility
 - » In-network aggregation
- **Medium-access control**

Peter A. Steenkiste, CMU

14

Simple Model for Energy Consumption



Typical values:

$$\alpha = 2 \dots 6$$


$$E_{elec} = 50 \text{ nJ/bit}$$

$$\epsilon_{amp} = 100 \text{ pJ/bit/m}^\alpha$$

Peter A. Steenkiste, CMU

15

Energy-efficient Routing : Example



Transmitting a single k -bit message from node A (located at distance nr from Sink) to Sink:

Direct transmission: $E_{direct} = E_{Tx}(k, d = nr) = E_{elec} * k + \epsilon_{amp} * k * (nr)^\alpha = k(E_{elec} + \epsilon_{amp} nr^\alpha)$

Multi-Hop Transmission: $E_{multi-hop} = n * E_{Tx}(k, d = r) + (n-1) * E_{Rx}(k)$
 $= n(E_{elec} * k + \epsilon_{amp} * k * r^\alpha) + (n-1) * E_{elec} * k = k((2n-1)E_{elec} + \epsilon_{amp} nr^\alpha)$

MultiHop routing requires less energy than direct communication if: $\frac{E_{elec}}{\epsilon_{amp}} < \frac{r^\alpha (n^{\alpha-1} - 1)}{2}$

Assuming $\alpha = 3, r = 10m$, we get $E_{multi-hop} < E_{direct}$ as soon as $n \geq 2$

Peter A. Steenkiste, CMU

16

Minimum Energy in a Wireless Network

- **Problem:** for an arbitrary configuration of nodes, find (in a fully distributed way) the minimum cost spanning tree to and from a given *sink* node
- **Assumptions**
 - » Each node knows precisely its own location (e.g., using GPS)
 - » The power decreases with distance according to a power law with a known and uniform exponent α
 - » Each node can communicate with another node located at an arbitrary distance
 - » Nodes do not move
 - » Slightly different power model

sending:	td^α
receiving:	c

- **Example:**

A	B	C
●	●	●

Power to send from A to C via B:

$$td_{AB}^\alpha + td_{BC}^\alpha + c$$

Peter A. Steenkiste, CMU

17

Relay region

Relay region:

$$R_{i \rightarrow r} \equiv \{(x, y) \mid P_{i \rightarrow r \rightarrow (x, y)} < P_{i \rightarrow (x, y)}\}$$

We can expand this to:


$$td_{i,r}^\alpha + td_{r,(x,y)}^\alpha + c < td_{i,(x,y)}^\alpha$$

$$t \left((i_x - x)^2 + (i_y - y)^2 \right)^{\alpha/2} - t \left((r_x - x)^2 + (r_y - y)^2 \right)^{\alpha/2} >$$

$$t \left((i_x - r_x)^2 + (i_y - r_y)^2 \right)^{\alpha/2} + c$$

RELAY
REGION

$R_{i \rightarrow r}$




Peter A. Steenkiste, CMU

18

Distributed Network Protocol

- Finds the minimum power topology for a stationary set of nodes with a single sink
- Assumption: each node is equipped with a GPS-receiver and transmits its position to its neighbors
- The protocol proceeds in 2 phases:
 1. Each node computes its own enclosure
 2. Each node computes its optimal cost distribution



Peter A. Steenkiste, CMU

19

Load-balancing


- Assumption: in a multi-hop many-to-one sensor network, the data collection follows a spanning tree.
- Power consumption due to transmission/reception grows exponentially from the leaves to the root of the tree
- Consequence: the power sources of the nodes close to the sink deplete faster. Since they relay all the network's traffic, they pull the network lifetime down.

Peter A. Steenkiste, CMU

20

Load-balancing


Line topology



P_{tr} : Average transmission power consumption
 P_{rc} : Average reception power consumption
 P_p : Average processing power consumption
 $P_t(k)$: Total power consumption of node k
 $P = P_{tr} + P_{rc} + (k-1)(P_p + P_{rc})$
 P grows linearly with the distance from the leaf node

Peter A. Steenkiste, CMU

Tree topology





d : distance from leaf
 F : number of messages forwarded
 P : Power consumption
 Assumptions:
 1) all nodes have either 0 or $n_c > 2$ children
 2) all leaves are at the same distance from the sink
 $F(d) \geq 2^d$
 $P(d) \geq P_{tr} + 2^d(P_p + P_{rc})$
 P grows exponentially with distance from leaf node

21

Load balancing

- Power consumption increases at least linearly when nodes are closer to the sink
- Typical case is much worse





Peter A. Steenkiste, CMU

22

Use Mobility for Load-balancing

- Move the base station to distribute the role of “hot spots” (i.e., nodes around the base station) over time
- The data collection continues through multi-hop routing wherever the base station is, so the solution does not sacrifice latency



Peter A. Steenkiste, CMU

23

In-network Data Aggregation

- To mitigate cost of forwarding, compute relevant statistics along the way: *mean*, *max*, *min*, *median* etc.
- Forwarding nodes aggregate the data they receive with their own and send one message instead of relaying an exponentially growing number of messages
- Issues
 - » Location-based information (which nodes sent what) is lost
 - » Distributed computation of statistics
 - *mean*: node needs to know both the mean values and the sizes of samples to aggregate correctly
 - *median*: only an approximated computation is possible
- Especially useful in a query-based data collection system
 - » Queries regard a known subset of nodes
 - » Aggregation function can be specified

Peter A. Steenkiste, CMU

24

Medium-Access Control

- **MAC attributes:**
 - » Collision avoidance
 - » Energy efficiency
 - » Scalability and adaptivity
- Nodes transmit very intermittently, but once a transmission is taking place, we must ensure that the intended receiver gets it.
- Current-consumption in receive state or in radio-on idle state are comparable
- Idle state (idle listening) is a dominant factor in power consumption


Goal is to put nodes to sleep most of the time, and wake them up only to receive a packet

Peter A. Steenkiste, CMU

25

Synchronous MACs

- **TDMA (similar to cellular networks)**




- **Shortcomings**
 - » Necessity to organize nodes in clusters and cluster hierarchies
 - » High control traffic cost
- **Possible solution**
 - » Each node maintains two schedules
 - Its parent schedule
 - The schedule it sets for its children
 - » Beacons are used to compensate for clock drifts

Peter A. Steenkiste, CMU

26

Asynchronous: B-MAC

- **Asynchronous**
- **Low Power listening**



Peter A. Steenkiste, CMU

27

Shortcomings

- Transmitting a packet is very expensive
- Overhearing is expensive
- Relaying packets is expensive (multihop)

Simple Improvement:

- **Aggregating packets before sending them**
 - » In low duty cycle data collection network, gain may be substantial
 - » Price to pay : real-time

Peter A. Steenkiste, CMU

28

Design Issue: Reliable data collection

- Many-to-one communication paradigm
- Multi-hop communication
- Nodes select one parent to send their data packets (tree topology)

Peter A. Steenkiste, CMU

29


MintRoute: A Data Collection Tree Routing Protocol

- Distance-vector routing protocol: one routing path per node
- A shortest path: Minimum number of transmissions
- The base station send periodic beacons that are broadcast by each node after incrementing a hop count
- Node select beacons with lowest hop count from the ones it received, and adds its sender among a list of potential parents
- Neighboring nodes exchange periodic beacons for link quality evaluation (gaps within the sequence # of packets → packet losses)
- Nodes select their parent based on hop count, link quality and load.
- Volatile routing topology → load balancing
- Cycle avoidance : Link quality must not vary too rapidly

Peter A. Steenkiste, CMU

30


MintRoute : Root Beacons



Peter A. Steenkiste, CMU

31

MintRoute : Link Estimation



Peter A. Steenkiste, CMU

32

Design Issue: Data Processing

- **Challenge:** develop a tool that makes it convenient to query any wireless sensor network
- **Cope with**
 - » Data heterogeneity
 - » Geographic distribution
 - » Platform diversity
 - » Etc...
- **Web-based interface to**
 - » Add/remove instances of sensor networks
 - » Display data
 - » Export data
 - » Send queries
- **Example:** <http://gsn.sourceforge.net/>



Peter A. Steenkiste, CMU

33

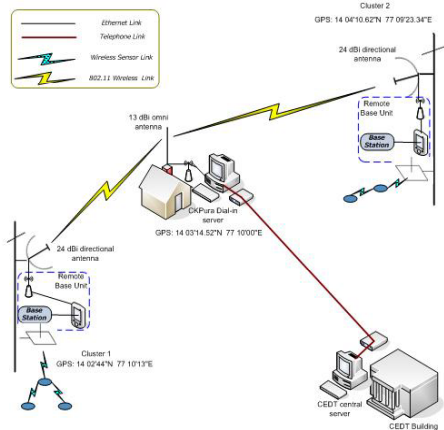
Design Issue: Hybrid Network Architecture

- Deployment can be in remote areas
- The sink needs to communicate the data to the user
- Pervasive computing means pervasiveness of access to data
- Multi-tier architectures are necessary

Peter A. Steenkiste, CMU

34

2-tier Architecture with 802.11 Bridge



Peter A. Steenkiste, CMU

35

Delay Tolerant Network with Data Mules

- Clusters are not directly connected to the server
- Cluster heads store data from the cluster nodes
- “Data mules” collect the data periodically
 - » Cars, robots, plane, etc.
- When a cluster-head detects a mule, it uploads to it the data it had in store

Peter A. Steenkiste, CMU

36

How Large is Too Large?

- **Largest Scale demonstration**
 - » ExScal : Ohio State University, DARPA-NEST project
 - Intrusion detection application (perimeter, pipeline, border, etc)
 - Demonstrated 3 tier WSN consisting of 1000 XSM motes, 200 stargate nodes, Dec 2004
 - 3rd tier controls stargates, each stargate controls 20-50 motes
 - 1.3 km x 300 m, 45 m spacing
 - Demo that lasted for a few days
- **Examples of long term POCs (double digit motes)**
 - » Habitat monitoring on Great Duck Island, UCB
 - » Golden gate bridge, UCB
 - » Condition based maintenance on BP oil tanker and in fab, Intel/Crossbow/Rockwell/BP
 - » Water Pipeline Monitoring, Intel
 - » Vineyard monitoring in Montepaldi farm, Good Food Project

Peter A. Steenkiste, CMU

Based on Lama Nachman's SECON 2006 panel talk

37

Condition-Based Management

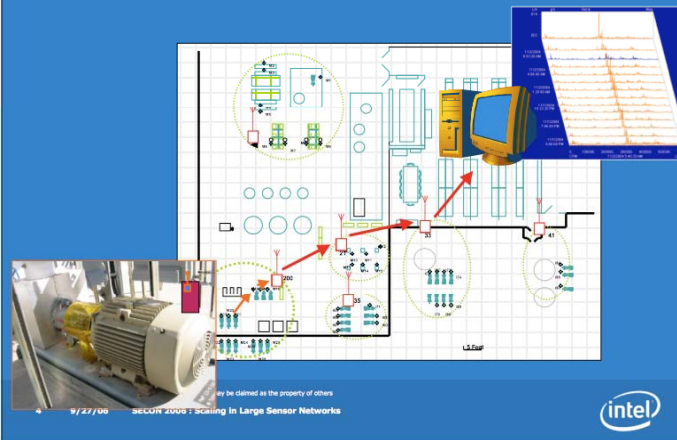
- **Use vibration signatures to identify problems with equipment**
- **Manual data collection common today**
 - » ~4000 sensor points in a fabrication plant
 - » ~200 KB of raw data per sensor point Lower operating costs
 - » Reduced manual labor
 - » Increased fault coverage
- **Prototype in Intel Fab and BP oil tanker**
 - » Compare data quality and reliability to manual system
 - » 25 node deployment
 - 6 Vibration/Temp channels per node
 - Tachometer



Peter A. Steenkiste, CMU

38

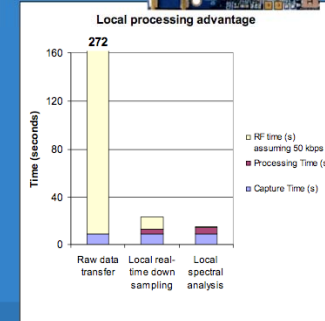


Condition-Based Maintenance Network



Peter A. Steenkiste, CMU

39

Advantage of edge processing



© Intel Corporation, 2006
 * Other names and brands may be claimed as the property of others
 SECON 2006 : Scaling In Large Sensor Networks

Peter A. Steenkiste, CMU

40

How Large can Sensor Networks Feasibly Grow?

- **Multi-tier is the typical approach to scaling**
 - » Multi-tier architecture divides the scaling problem into manageable segments
 - » Segmented networks (2nd tier aggregation, high throughput apps)
 - Static allocation of networks simplifies the scalability problem but complicates the management problem
 - Dynamic allocation of networks is more appealing
- **Clusters on the orders of 10-50 nodes are adequate for most applications**
 - » In most industrial applications, having access to power for 2nd tier nodes is realistic (or at least intermittent power sufficient to charge a battery)
 - » Outdoor deployments can also leverage solar power for recharging 2nd tier batteries
 - » Pipeline monitoring applications look different, will probably need higher scaling

Peter A. Steenkiste, CMU

41

Conclusion Sensor Networks

- **WSNs are an emerging technology which is poised to grow exponentially in the coming years**
- **This new communication paradigm introduces a new set of design constraints**
 - » They must be extremely low-cost
 - Both to purchase and to operate
 - » They must be extremely energy efficient since their lifetime is potentially years
 - Hardware design
 - Routing and topology mechanisms
 - Specialized Medium Access Control mechanisms
 - » Despite their low-cost and power management features, they must implement reliable communication protocols
 - » They must integrate versatile middle-ware and provide data processing
 - » They will often rely on a hybrid network infrastructure

Peter A. Steenkiste, CMU

42