Mobile and Ubiquitous Computing
Sensor network applications and research challenges

Niki Trigoni
www.dcs.bbk.ac.uk/~niki
niki@dcs.bbk.ac.uk

Overview

- An overview of sensor networks
- Application scenarios and requirements
  - habitat monitoring
  - traffic control
  - emergencies
- Research problems
  - network deployment and configuration
  - query processing and storage management
  - network longevity and robustness
  - other research issues
- Discussion

Sensor node

Node with

- computation,
- storage,
- communication and
- sensing capabilities.

Example: Tmote

- 8 MHz Processor
- 10k RAM, 48k Flash
- 250kbps 2.4GHz Radio
  - 50m range indoors / 125m range outdoors
- Integrated Humidity, Temperature, and Light sensors
- Programming and data collection via USB
- TinyOS support

Sensor network

- A collection of sensor nodes deployed in an area and connected through a multi-hop wireless network.

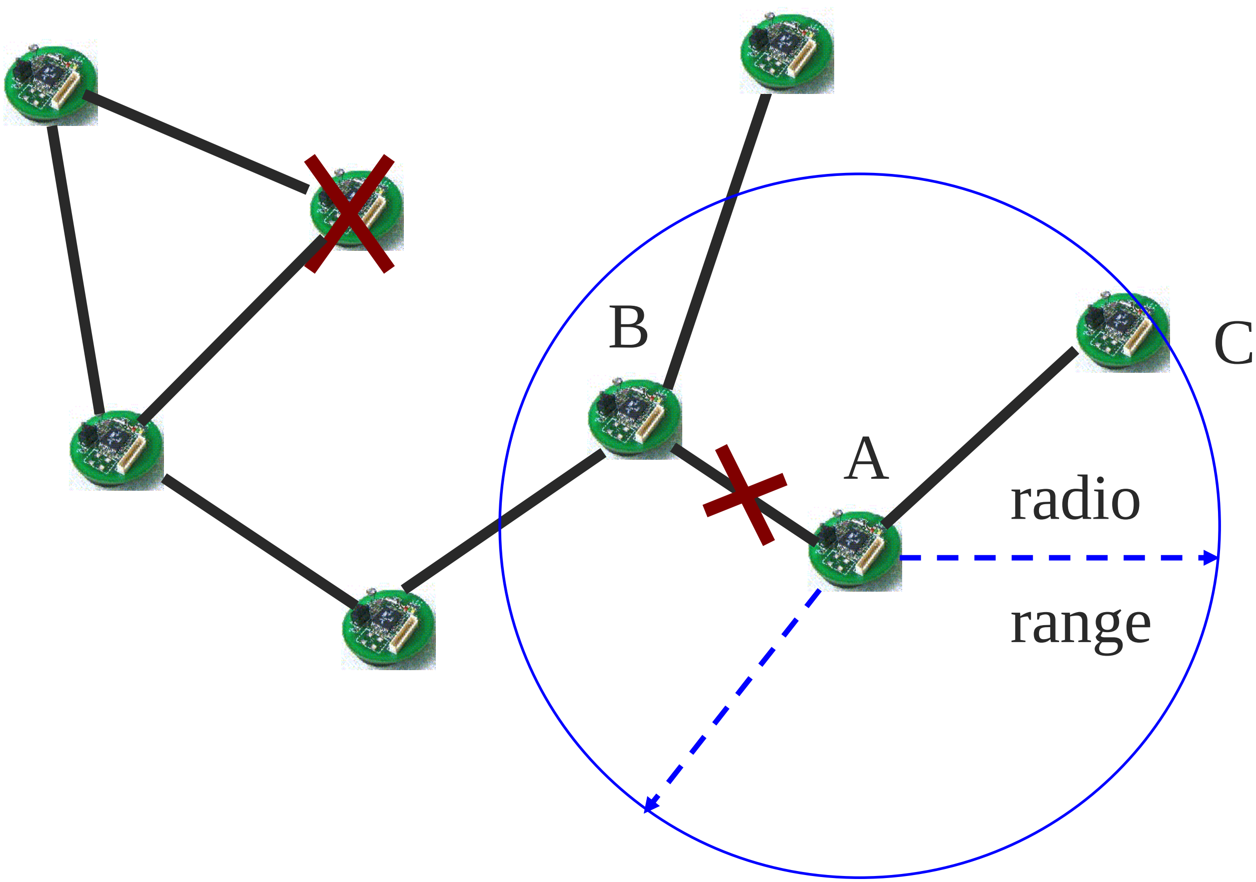

Simple deployment

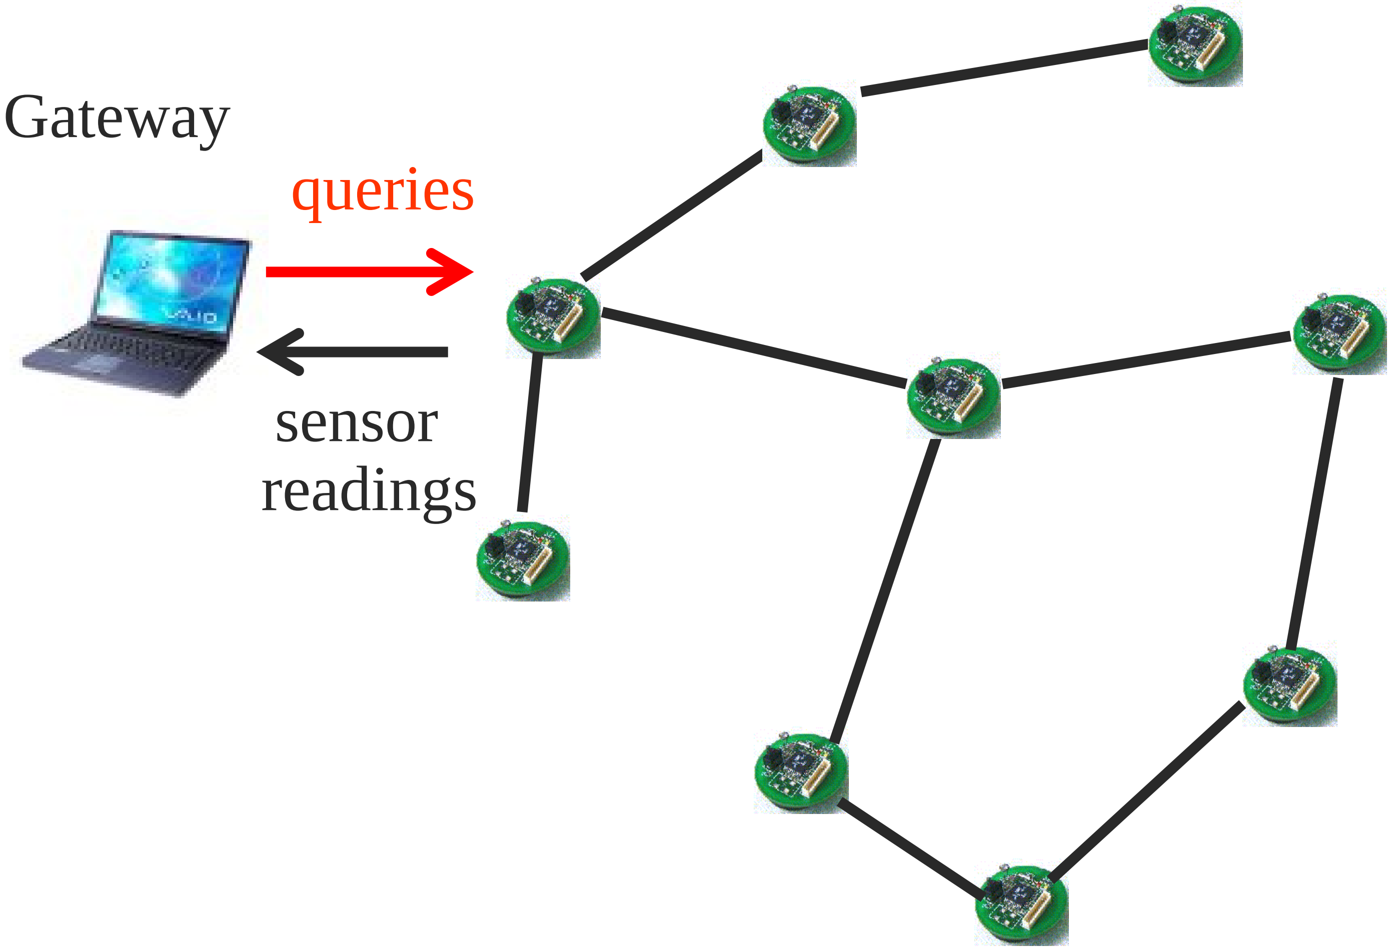

Sensor node platforms

Node	Bandwidth (Kbps)	MIPS	Flash (Mb)	RAM (Kb)	Duty Cycle %
Spec	50	5	0.1	4	0.1-0.5
Mote	100	10	0.5	10	1-2
Imote	500	50	10	128	5-10
Stargate	10000	100	32	512	50-100

J. Hill, M. Horton and R. Kling (ACM Comm. June 2004)

New platforms have emerged since 2004 …

Hierarchical deployment

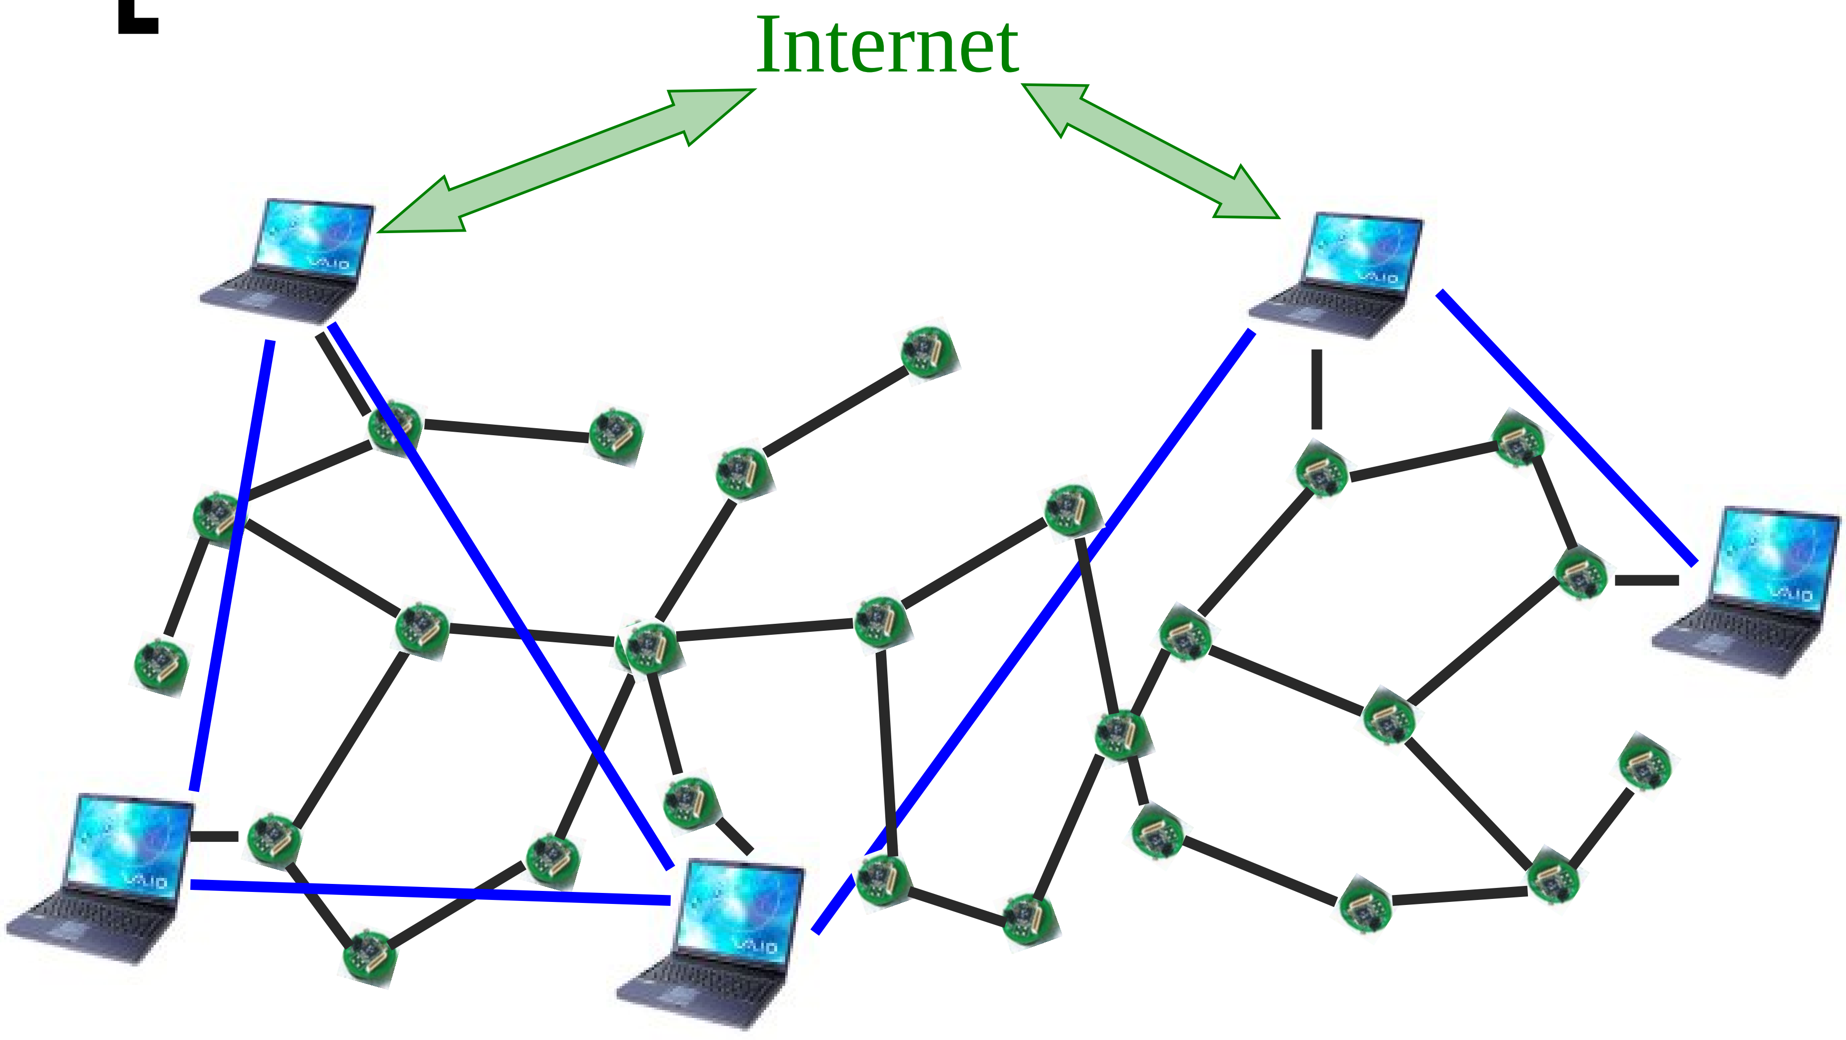

Other hybrid deployments

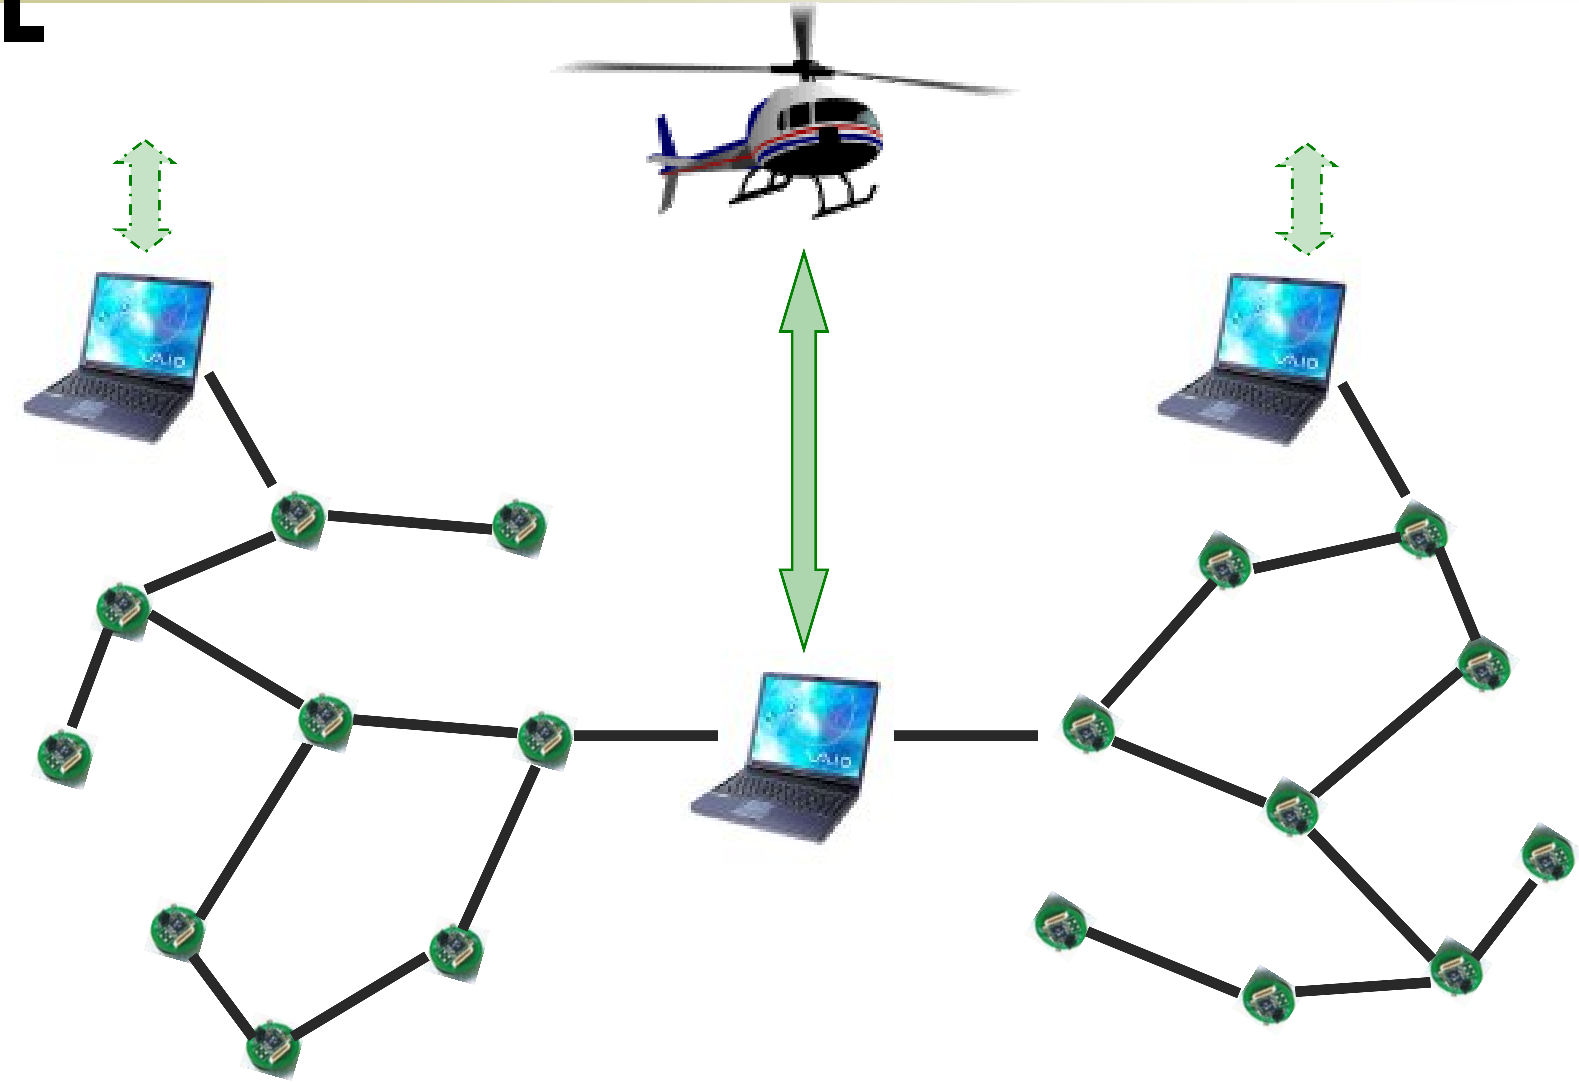

The DB view of sensor networks

Traditional

Procedural addressing of individual sensor nodes; user specifies how task is executed; data is processed centrally.

DB Approach

Declarative querying; user isolated from “how the network works”; in-network distributed processing.

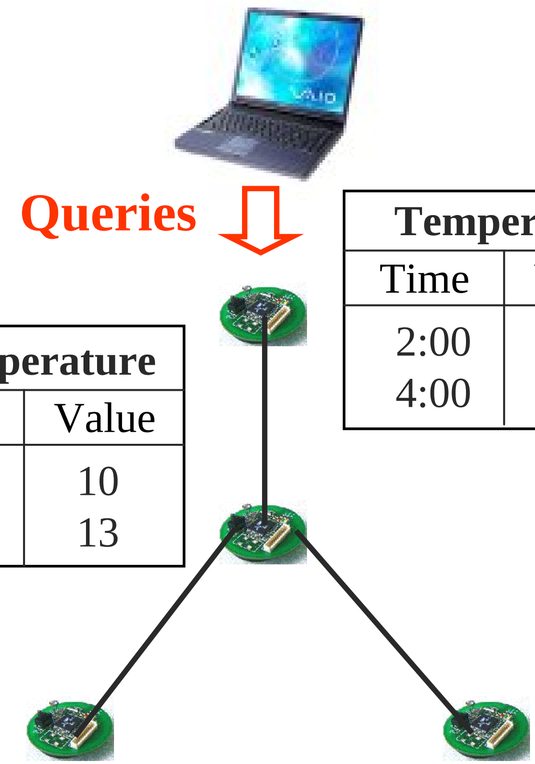

Temperature	
Time	Value
2:00	15
4:00	12

Temperature	
Time	Value
2:00	10
4:00	13

Humidity	
Time	Value
2:30	70
3:30	75

Temperature	
Time	Value
2:00	20
3:00	18

Pressure	
Time	Value
1:00	30
4:00	35

Query examples

Snapshot queries:

- How many empty bird nests are in the northeastern quadrant of the forest?

```
SELECT      SUM(s)
FROM        SensorData s
WHERE       s.nest = empty and  s.loc in (50,50,100,100)
```

Long-running queries:

- Notify me over the next hour whenever the number of empty nests in an area exceeds a threshold.

```
SELECT      s.area, SUM(s)
FROM        SensorData s
WHERE       s.nest = empty
GROUP BY    s.area
HAVING      SUM(s) > T
DURATION    (now, now+60)
EVERY       5
```

Sensor network applications

Examples:

- habitat monitoring
- chemical and biological sensors
- fire, earthquake emergencies
- vehicle tracking, traffic control
- surveillance of city districts
- defense-related networks
- alerts to terrorist threats
- …

Habitat monitoring (HM)

Sensor networks in a national park:

- capture micro-climates
- monitor animal behavior
- cover large areas over large periods of time
- adjust sensing devices to a suitable degree of precision
- identify changes in the habitat
- report unusual or seasonal events, like bird migration
- identify emergencies, like fires and prevent their expansion
- set alarms for contaminated land or water areas
- sense approaching visitors and provide online directions
- help in properly maintaining the park’s infrastructure etc.

Examples of HM systems

- Great Duck Island Project
  - Intel Research Lab at Berkeley
  - College of the Atlantic in Bar Harbor
  - University of California at Berkeley

  2002: over 1 million readings logged from 32 Mica motes
  2003: 150 nodes, 25 weather station nodes

- James Reserve Extensible Sensing System
  - University of California (CENS)

  microclimates and animal detection within a 25-hectare area
  Mica motes and Compaq iPAQs deployed in over 100 locations

- PODS in Hawaii Volcanoes National Park
  - University of Hawaii

  monitor micro-climates in areas with endangered species of plants

Examples of HM systems

- CORIE
  - Oregon Graduate Institute

  13 stations located throughout the Columbia river estuary

  measure field velocity, salinity, temperature and water level

- ORCA
  - University of Washington
  - Washington State Department of Ecology

  an autonomous water quality monitoring system part of the U.S. Coastal Intensive Sites Network (CISNet)

- Floodnet, Secoas, GlacsWeb
  - University of Southampton and partners

  monitor weather, sea bed movement and sub-glacial movement resp.

Traffic surveillance and control

Sensor networks can assist in several ways:

- detect traffic hotspots and warn approaching drivers
- divert traffic and increase transportation capacity
- monitor roads for accidents and car failures
- manage parking spaces
- detect illegal driving and parking behavior
- monitor continuously the condition of roads and signal repairs
- road tolling

Traffic surveillance and control

Examples:

- Traffic Monitoring Sensor Network

  UC Berkeley, MLB Company

  29 palms deployed from an unmanned aerial vehicle (UAV)

  measure vehicle velocity and closest point of approach

- Traffic monitoring in the city of Cambridge

  Univ. of Cambridge, Birkbeck College

  Deploy sensor nodes to monitor incoming and outgoing traffic at various intersections/roundabouts, process and store data and provide high-level traffic information to interested parties

Emergency scenarios

Sensor networks can assist in several ways:

- identify early signs of fire in forests
- help fire fighters predict the direction in which fire expands
- prevent fire fighters from getting trapped
- alert people about imminent flood
- assist in rescue operations, e.g. by locating victims or members of the rescue team
- identify fire, flood and gas leaks in smart buildings, activate reactors, isolate rooms, help people evaluate the building

Application requirements

General requirements

- self-configurable network
- wireless multi-hop communication of sensor readings
- querying ability

Habitat monitoring

- communication from multiple data sources to a sink
- large-scale deployment
- unattended operation for long periods of time

Traffic control

- collaborative operation (track a vehicle, predict bottleneck)

Emergencies

- robustness to node and link failures
- bursty traffic vs. regular traffic

Roadmap

- An overview of sensor networks
- Application scenarios and requirements
  - habitat monitoring
  - traffic control
  - emergencies
- **Research problems**
  - **network deployment and configuration**
  - **query processing and storage management**
  - **network longevity and robustness**
  - **other research issues**
- Discussion

Network deployment and configuration

Localization:

- How do nodes (without GPS) infer their position?
- *Distributed localization algorithms* enable nodes to estimate their position after interacting with their peers.
  - determine the distance from anchor nodes
  - infer node position from those distances
  - refine position estimate based on distance from neighbors, and their position estimates

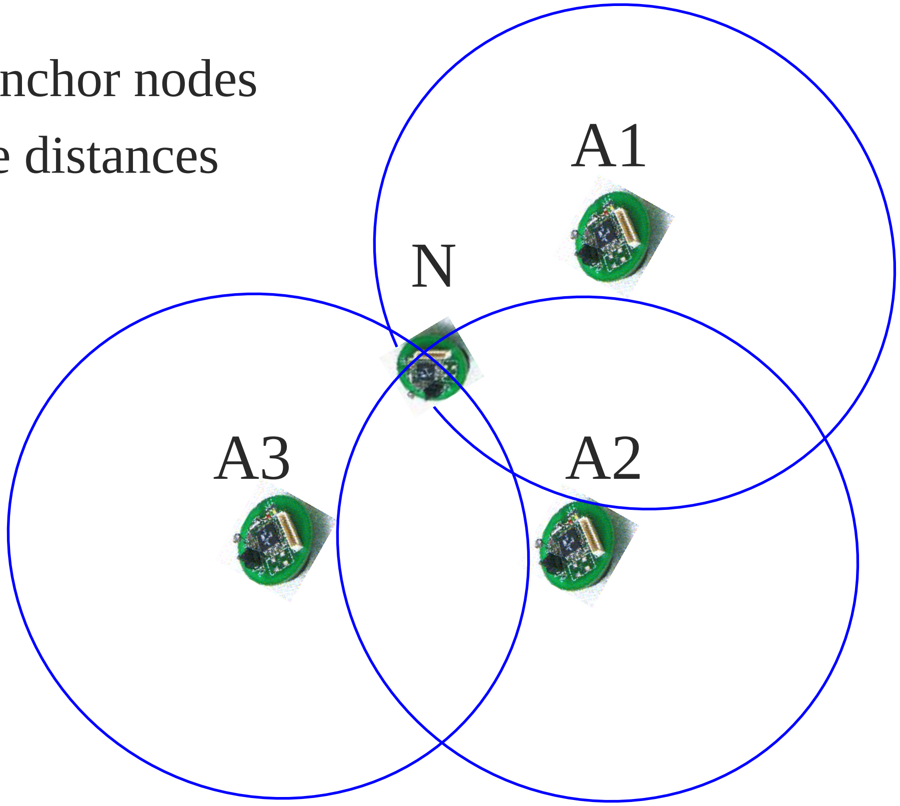

Doherty et al., Estrin et al., Li et al., Langendoen et al., Priyantha et al., Savvides et al., Whitehouse et al., …

Network deployment and configuration

Time synchronization:

- Sensor readings must be annotated with timestamps.

  How do nodes keep track of time?

- *Distributed algorithms* (like TPSN and RBS) synchronize neighbor nodes to a global time or relatively to each other.

  BUT:

  - Synchronization error (drift) increases with the distance from the global clock.
  - Perfect synchronization is impossible.
  - Need for scalable solution

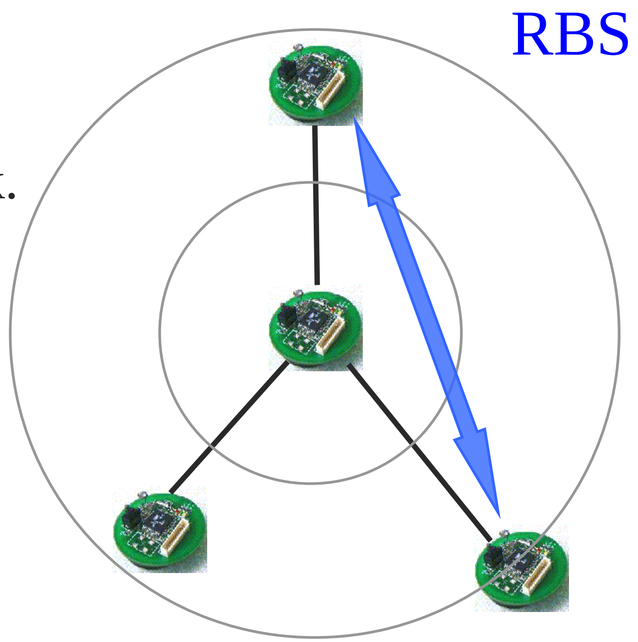

Elson et al.(RBS), Gareniwal et al.(TPSN), Girod et al., Lamport, Mills, Römer

Network deployment and configuration

Time synchronization:

- Sensor readings must be annotated with timestamps.

  How do nodes keep track of time?
- *Distributed algorithms* (like TPSN and RBS) synchronize neighbor nodes to a global time or relatively to each other.

  BUT:
  - Synchronization error (drift) increases with the distance from the global clock.
  - Perfect synchronization is impossible.
  - Need for scalable solution

TPSN

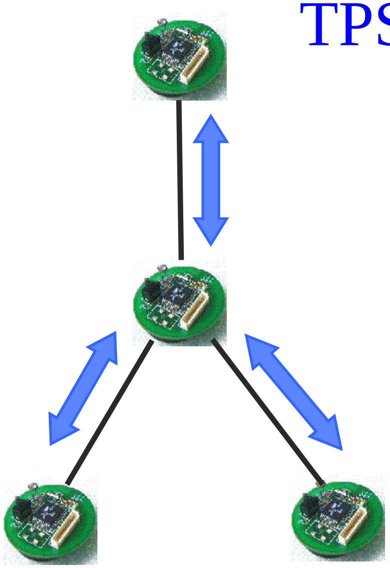

Elson et al.(RBS), Gareniwal et al.(TPSN),

Girod et al., Lamport, Mills, Römer

Query processing and routing

Processing aggregate queries:

- Centralized processing
- ✓ In-network processing

    Madden et al (TAG),

    Chalermek et al. (Directed diffusion),

    Gehrke et al.(Cougar)

With centralized processing, results at each edge grow linearly in the number of descendant nodes.

```
SELECT  SUM(s)
FROM    SensorData s
WHERE   s.nest = empty
EVERY   60 min
```

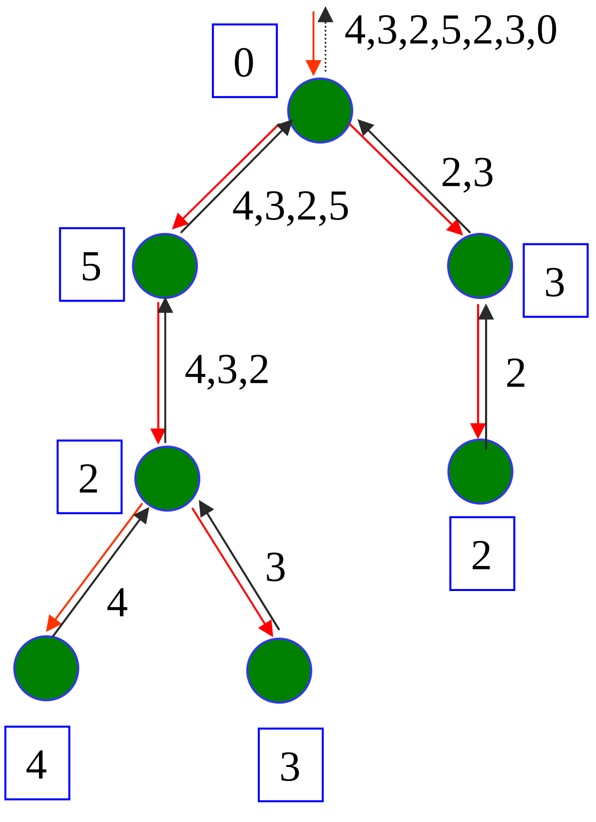

Query processing and routing

Processing aggregate queries:

- Centralized processing
- ✓ In-network processing

    Madden et al (TAG),

    Chalermek et al. (Directed diffusion),

    Gehrke et al.(Cougar)

With in-network processing, the number of results at each edge remains constant.

⇒ Reduces communication overhead

⇒ Reduces energy consumption

⇒ Increases network lifetime

```
SELECT   SUM(s)
FROM     SensorData s
WHERE    s.nest = empty
EVERY    60 min
```

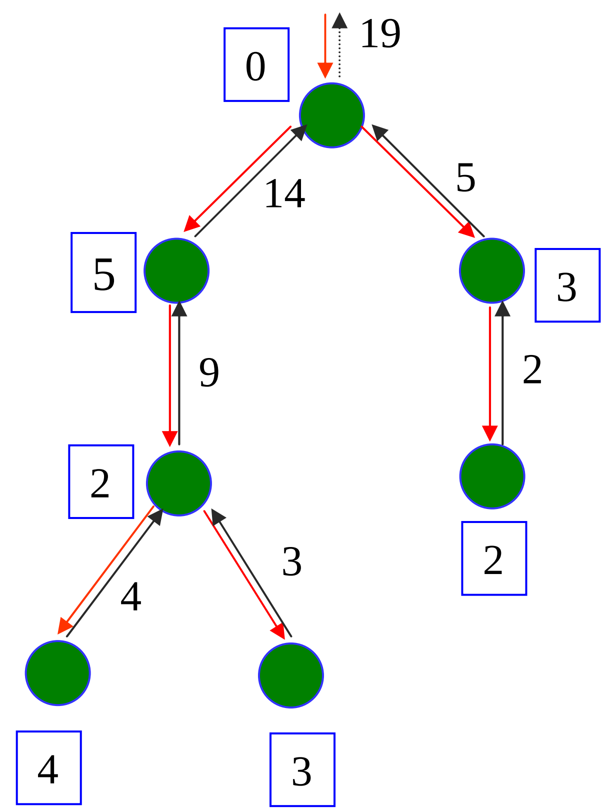

Probabilistic sensor updates

What if not all sensors generate new readings at regular intervals?

- return results of updated queries ?
- return values of updated sensors ?

Query and storage management

Hybrid pull-push data dissemination:

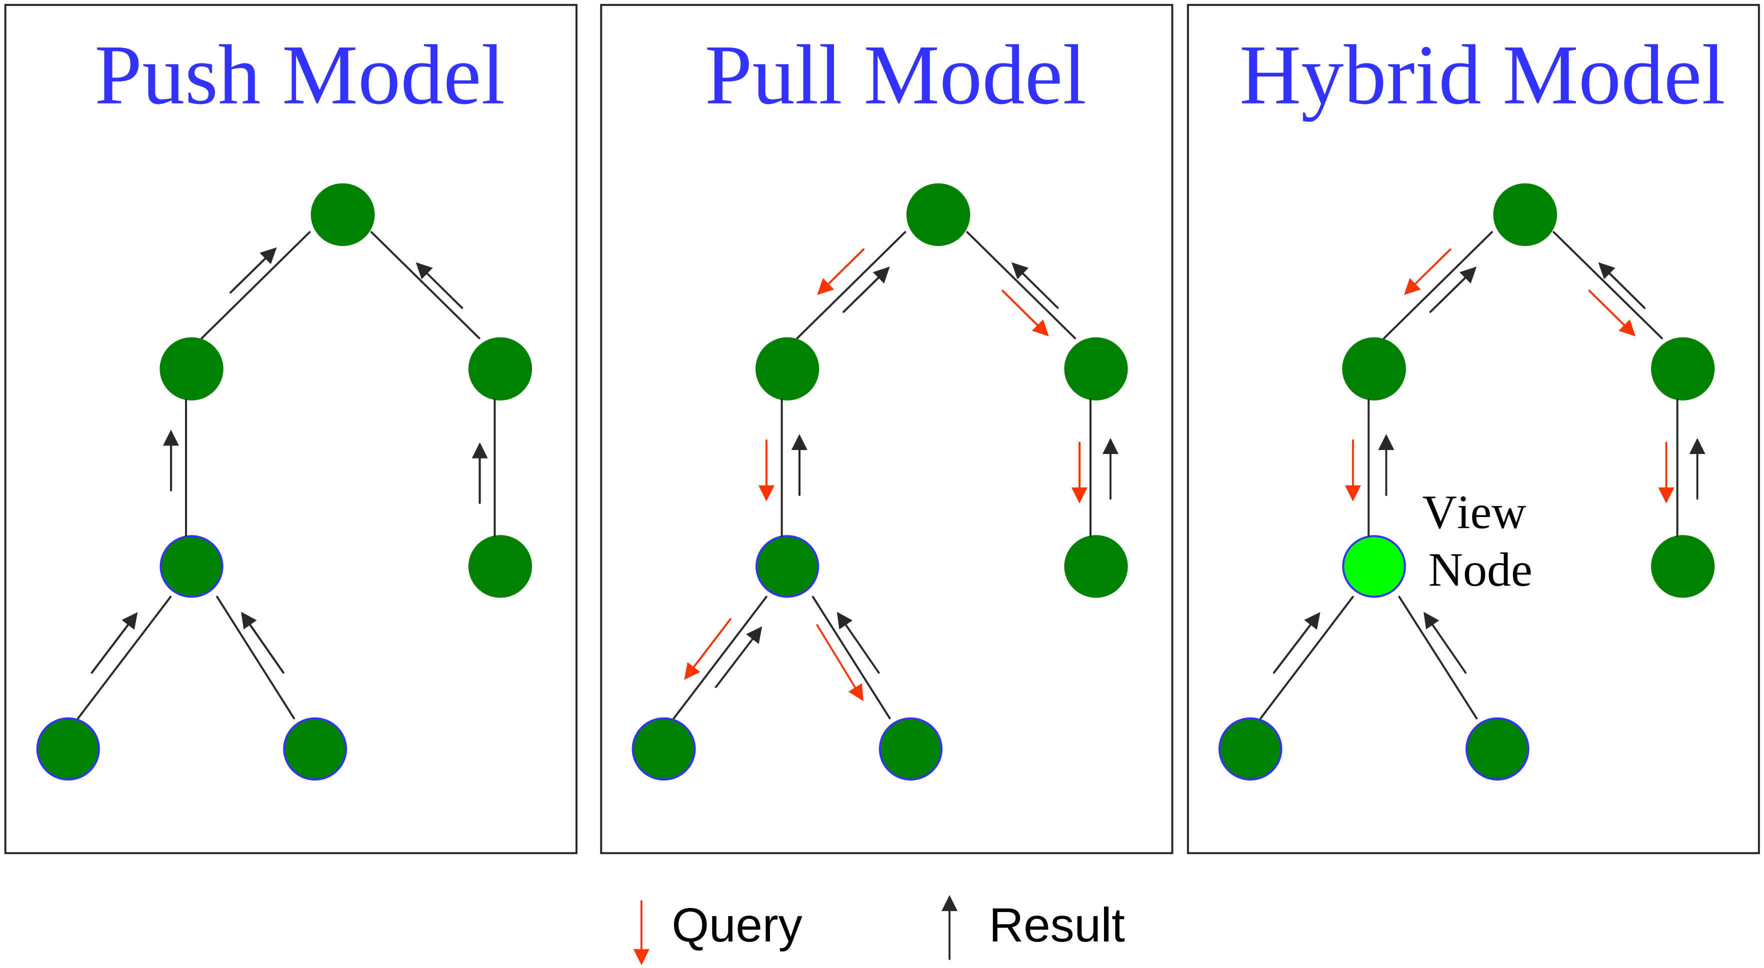

Storage management

Data-centric storage:

- GHT (Geographic Hash Table) hashes sensor readings to locations where they should be stored (like in P2P systems).
- It then stores the data at sensor nodes close to these locations.
- Queries retrieve sensor readings from the designated storage nodes.
  - How is the hash function selected?
  - What kind of data should be stored (raw readings vs. processed data)?

Ratnasamy, Shenker, Karp, Govindan, Estrin, Yin, Yu

Extending the lifetime of the network

- Reducing the number of transmitted messages is not the only way of saving energy at the nodes.
- Recent studies about radio energy consumption:

listen	*receive*	*transmit*	
1	1.05	1.4	[Stemm et al. 97]
1	2	2.5	[Kasten 01]
1	1.2	1.7	[Chen et al. 02]
1	1.1	3	(Mica motes)

Listening is *not* much cheaper than sending/receiving.

- Nodes should try to turn off the radio for as long as possible, in order to extend the lifetime of the network.

Extending the lifetime of the network

Energy-aware participation in routing and sensing operations:

- Topology control schemes (e.g. GAF)
- As node density increases, there is redundancy in:
  - nodes participating in routing backbone
  - number of sensing devices

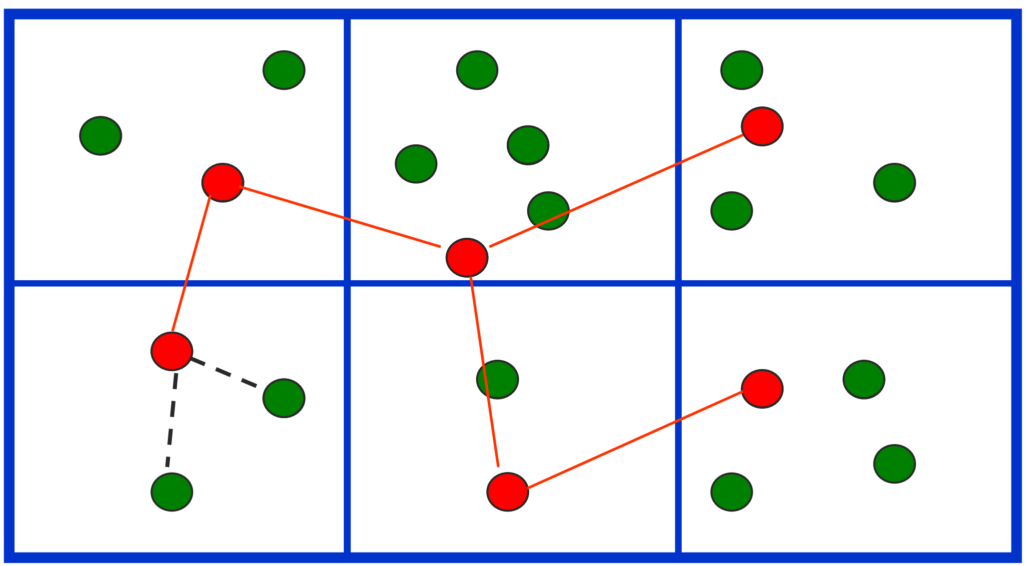

e.g. GAF [Xu et al.]:
Geographic
Adaptive
Fidelity

Extending the lifetime of the network

Node scheduling:

- Even those nodes that decide to participate in routing do not need to keep their radios on all the time
- Ideally radio is on only when sending or receiving data
- A few examples:
  - Tree-based scheduling (Madden et al.)
  - S-MAC (Heidemann et al.)
  - …
- Tradeoff between energy consumption at the nodes and delay in message delivery

Extending the lifetime of the network

TAG Node Scheduling [Madden et al. 2002]

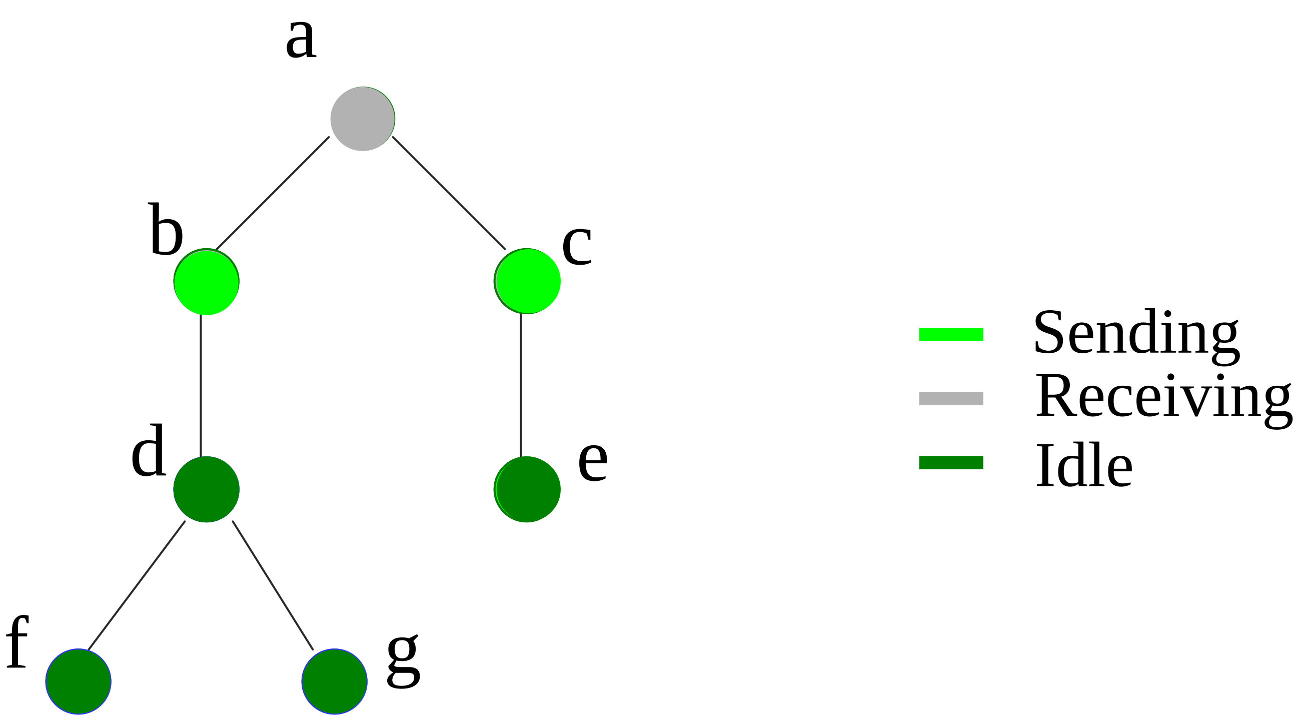

Robustness

Resilience to node/link failures and packet drops:

- route data thru multiple paths: node-disjoint vs. braided paths
- localized algorithms for the construction of alternate paths
- tradeoff between fault-tolerance guarantees and energy needed to ensure them
- in-network processing in the context of multi-path routing
- storage nodes: replication schemes that avoid data loss in case of node failures
- tradeoff between the resilience of different replication schemes and their maintenance overhead
- replication algorithms should take into account node capabilities

Other research issues

- Query processing
  - approximate queries
  - event-based queries
  - application-specific query requirements
- Security issues
  - resilience to malicious attacks
  - continuous network connectivity
  - access control
  - data integrity
- Privacy issues
  - transparency of recorded sensor information
  - legitimate use of sensed data
  - sensing within allowed areas

Discussion

- Sensor networks have the potential of assisting in many aspects of our life.
- Deploying and operating a large sensor network for long periods of time is not trivial.
- Energy-efficiency, fault-tolerance, security and privacy are important requirements for most sensor applications.
- These requirements should be taken into consideration in designing self-configurable sensor networks with query processing and storage capabilities.

Related reading

• Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. *TAG: a Tiny AGgregation service for ad-hoc sensor networks*. 5th Symposium on Operating System Design and Implementation, 2002.
http://db.cs.berkeley.edu/papers/osdi02-tag.pdf

To prepare for discussion:

• Ya Xu, John Heidemann, and Deborah Estrin. *Geography-informed Energy Conservation for Ad Hoc Routing*. In Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking, pp. 70-84, 2001.
http://www.isi.edu/~johnh/PAPERS/Xu01a.pdf