Part 1 --Sensor Network Applications

Peter Scheuermann & Goce Trajcevski

Military Applications:

Command, Control, Communications, Computing, Intelligence, Surveillance, Reconnaissance, Targeting (C4ISRT)

- Monitoring friendly forces, equipment and ammunition
- Battlefield surveillance
- Reconnaissance of opposing forces and terrain
- Targeting
- Battle damage assessment
- Nuclear, Biological and Chemical (NBC) attack detection and reconnaissance

Further Military Applications

- Intrusion detection (mine fields)
- Detection of firing gun (small arms) location
- Chemical (biological) attack detection
- Targeting and target tracking systems
- Enhanced navigation systems
- Battle damage assessment system
- Enhanced logistics systems

Environmental Applications

- Tracking the movements of birds, small animals, and insects
- Monitoring environmental conditions that affect crops and livestock
- Irrigation
- Earth monitoring and planetary exploration
- Chemical/biological detection
- Biological, Earth, and environmental monitoring in marine, soil, and atmospheric contexts
- Meteorological or geophysical research
- Pollution study
- Precision agriculture
- Biocomplexity mapping of the environment
- Flood detection, and Forest fire detection.

Habitat Monitoring

http://www.greatduckisland.net Great Duck Island in Maine

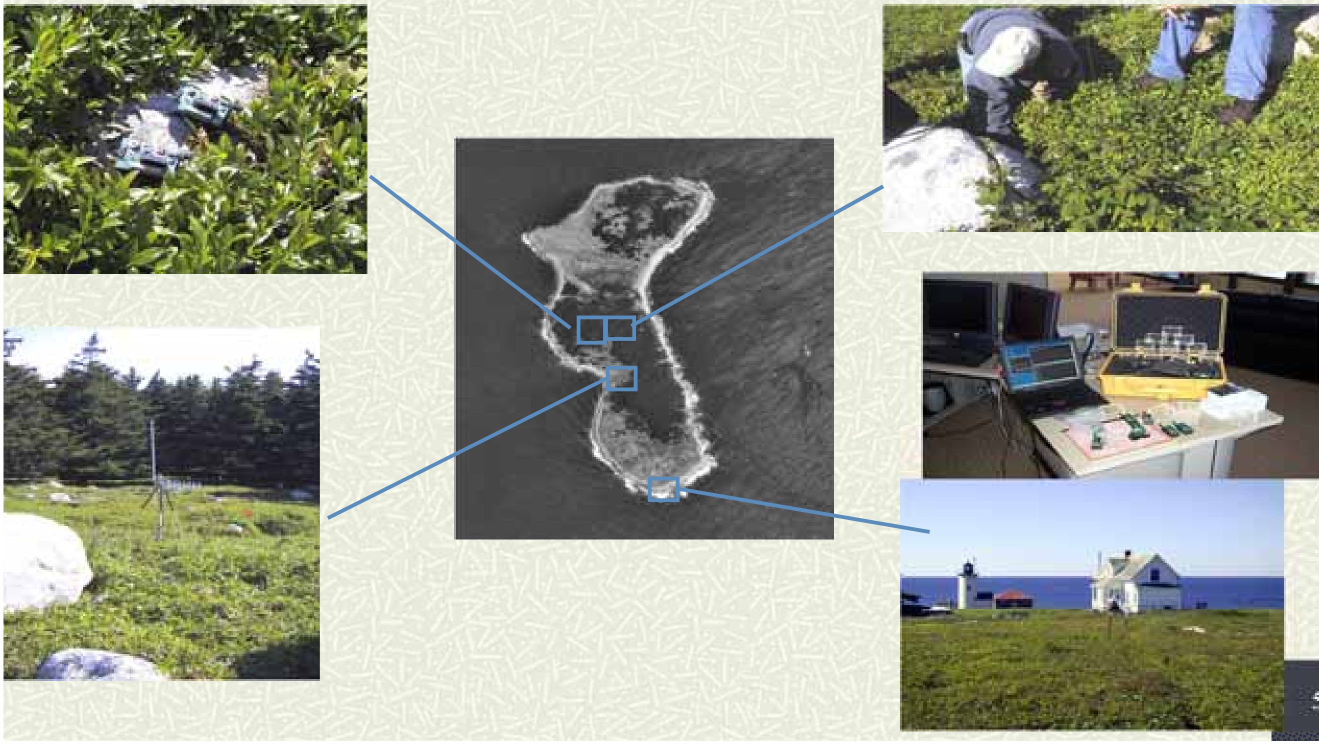

Motivation

Questions

- What **environmental factors** make for a good nest?
- How much can they vary?
- What are the **occupancy patterns** during incubation?
- What **environmental changes** occurs in the burrows and their surroundings during the breeding season?

Motivation--continued

Problems

- **Seabird colonies are very sensitive to disturbances**
- **The impact of human presence can distort results by changing behavioral patterns and destroy sensitive populations**
- **Repeated disturbance will lead to abandonment of the colony**

Solution

- **Deployment of a sensor network**

Habitat Monitoring

- Approx. 200 nodes including MICA, MICA2, burrow nodes (with IR) and weat station nodes
- Motes detect light, barometric pressure, relative humidity and temperature conditions.
- An infrared heat sensor detects whether the nest is occupied by a seabird, and whether the bird has company.
- Motes within the burrows send readings out to a single gateway sensor above ground, which then wirelessly relays collected information to a laptop compute lighthouse (~350 feet).
- The laptop, also powered by photovoltaic cells, connects to the Internet via sate
- Computer at base-station logs data and maintains database

GDI Sensor Network

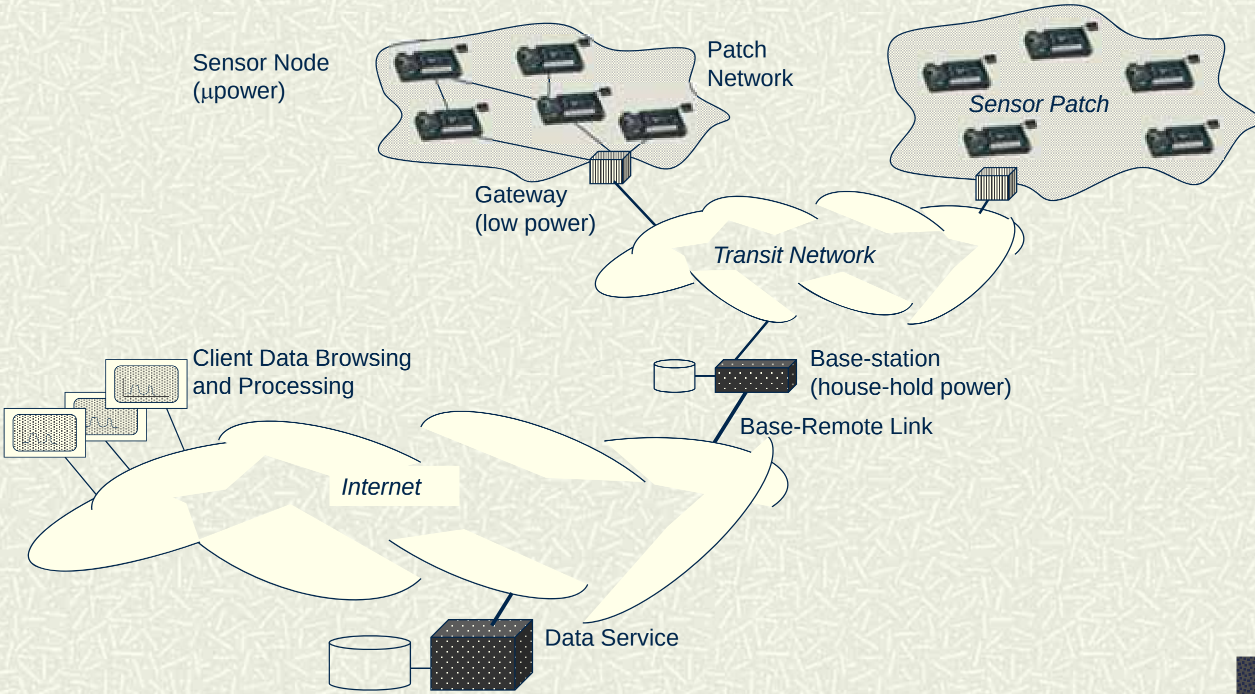

Communication

Routing

- **Routing directly from node to gateway not possible**
- **Approach proposed for scheduled communication:**
  - **Determine routing tree**
  - **Each gate is assigned a level based on the tree**
  - **Each level transmits to the next and returns to sleep**
  - **Process continues until all level have completed transmission**
  - **The entire network returns to sleep mode**
  - **The process repeats itself at a specified point in the future**

Network Re-tasking

Initially collect absolute temperature readings

- After initial interpretation, could be realized that information of interest is contained in significant temperature changes
- Full reprogramming process is costly:
  - Transmission of 10 kbit of data
  - Reprogramming application: 2 minutes @ 10 mA
  - Equals one complete days energy
- Virtual Machine based retasking:
  - Only small parts of the code needs to be changed

Health and Status Monitoring

- Monitor the mote’s health and the health of neighboring motes
- Duty cycle can be dynamically adjusted to alter lifetime
- Periodically include battery voltage level with sensor readings (0~3.3volts)
- Can be used to infer the validity of the mote’s sensor readings

IWESTERN
IVERSITY

Forest Fire Detection: FireWxNet
(Best paper award at Mobisys06—Hartung et al)

- Design and construction of a multi-tiered portable wireless system for monitoring weather conditions in rugged fire environments.
- Blends together long-range wireless technology based on directional radios with a short-range multi-hop sensor network.
- Also integrated web-enabled surveillance cameras to provide visual data

IWESTERN
IVERSITY

Weather and Forest Fires

mperature
- As Temperature goes up: fuels dry out, fuels temp rises, etc..

lative Humidity
- 10 minute fuels (small twigs) change with RH quickly.
- 10,000 hour fuels (large trees) not very affected.

nd Speed and Direction
- Feeds fire with extra oxygen. Dictates direction. Can change quickly.

ample: Temperature Inversions and ermal Belts
- Temperature is colder at lower elevations.
- Fires above the inversion continue to burn actively.
- Band of warmer air trapped midway up the mountain.
- When they break – increased winds, rapid increase in temp.

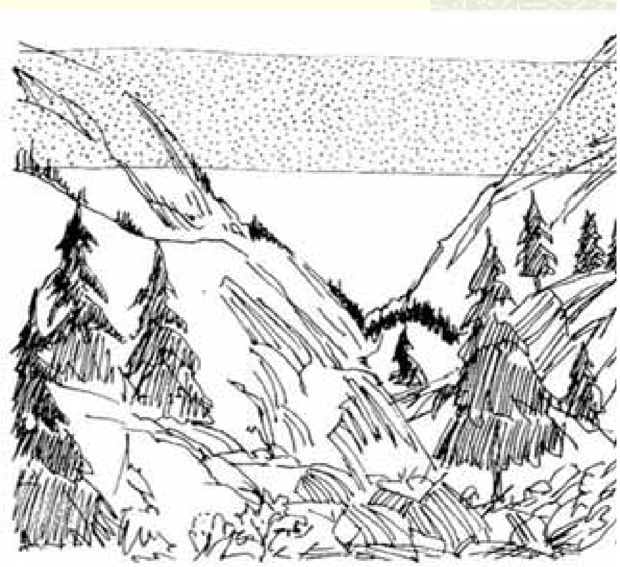

Forest Fires – Current Ways to Measure Weather

Belt Weather Kit

Disadvantages: poor coverage, often forgotten, slow, etc.

FireWxNet Design Goals

- **Deploy a WSN to monitor weather data around *prescribed* forest fires**
  - Temp, RH, wind
  - Actually deployed around *unplanned* lightning-strike wildfires
  - In collaboration with the University of Montana
- **Must be rugged, on-demand/portable, inexpensive, long-lived, easy to deploy, etc.**
  - Did not focus on survivability of nodes, but system is robust to node failures
- **Deployment difficulties => sparse WSN**
  - Focus on improved coverage of the thermal inversion belt – *well-known but difficult to detect*
  - capturing *vertical* elevational temperature differences, not horizontal, which would require 100's of sensor nodes

CU's Carl Hartung deploys WSNs in the fire area via helicopter

FireWxNet's Hierarchical Architecture

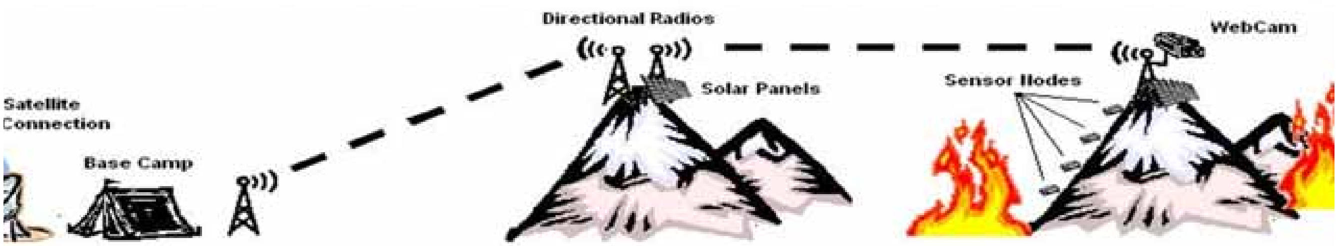

Soekris Base Station

FireWxNet's Wireless Sensor Network

- MICA2, 900 MHz, 2 AA batteries
- Mantis OS (MOS)
- Commercial sensors
  - e.g. Relative Humidity: Humirel 1520, extremely accurate below 30% RH

Networking Design:
Deployment, Duty Cycling, and Routing

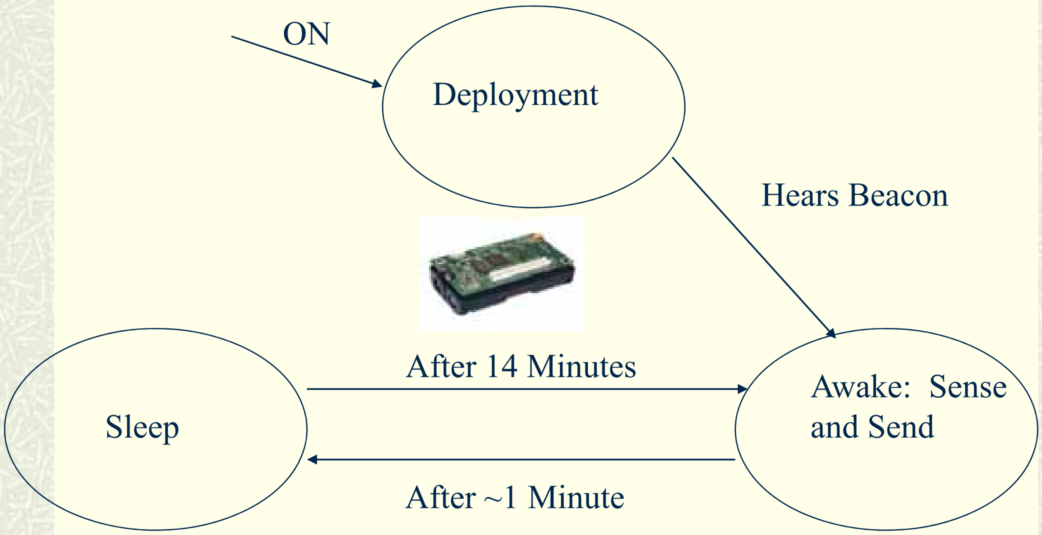

6% Duty Cycle – maintained by the beacons

Ease of Deployment

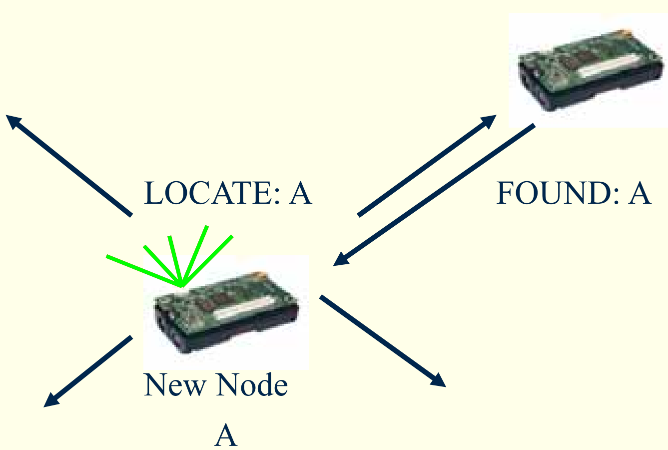

- Used to set up symmetric links where possible
- Crude but effective, but could be improved

Sensor Network – Routing

During 1 min Awake period,

- Beacons flooded every 4 seconds from Base
- Establishes tree-structured routing
- Data was sent 1 packet per second for 60 seconds

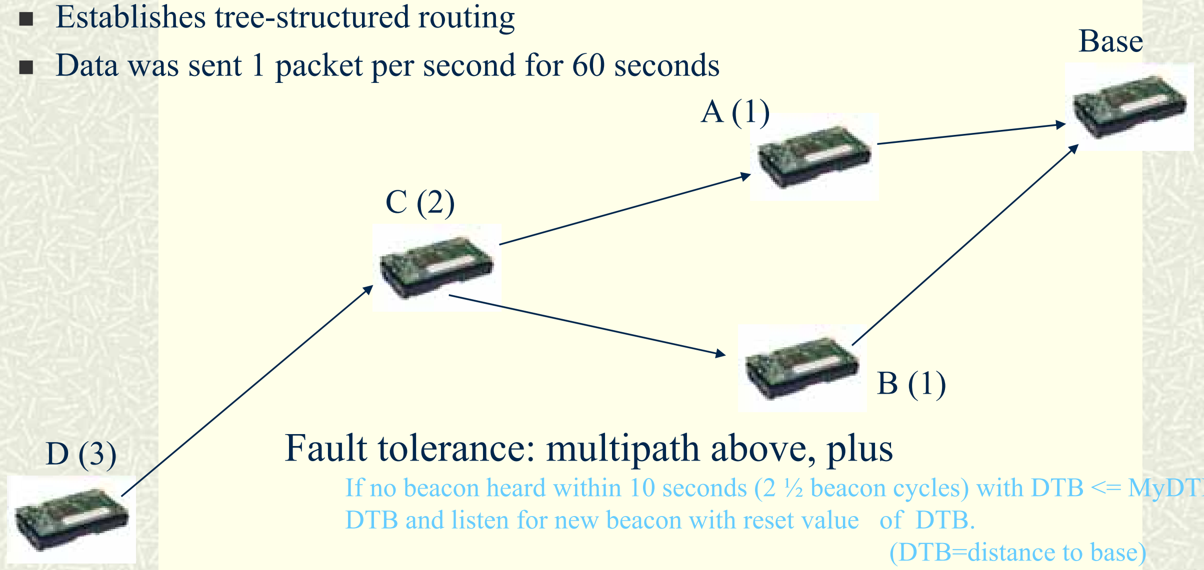

Fault tolerance: multipath above, plus

If no beacon heard within 10 seconds (2 ½ beacon cycles) with DTB <= MyDTB -1, reset DTB and listen for new beacon with reset value of DTB.

(DTB=distance to base)

Adapts to failed, moved, shut off nodes

Health Applications

- **Providing interfaces for the disabled**
- **Integrated patient monitoring**
- **Diagnostics**
- **Telemonitoring of human physiological data**
- **Tracking and monitoring doctors and patients inside a hospital, and**
- **Drug administration in hospitals**

CodeBlue: WSNs for Medical Care

http://www.eecs.harvard.edu/~mdw/proj/codeblue/

- NSF, NIH, U.S. Army, Sun Microsystems and Microsoft Corporation
- Motivation - Vital sign data poorly integrated with pre-hospital and hospital-based patient care records

CodeBlue: WSNs for Medical Care

are
- Small wearable sensors
- Wireless pulse oximeter / 2-lead EKG
- Based on the Mica2, MicaZ, and Telos sensor node platforms
- Custom sensor board with pulse oximeter or EKG circuitry
- Pluto mote
  - scaled-down version of the Telos
  - rechargeable Li-ion battery
  - small USB connector
  - 3-axis accelerometer

CodeBlue: WSNs for Medical Care

CodeBlue - scalable software infrastructure for wireless medical devices
- Routing, Naming, Discovery, and Security
- MoteTrack - tracking the location of individual patient devices indoors and outdoors

Heart rate (HR), oxygen saturation (SpO2), EKG data monitored

Relayed over a short-range (100m)

Receiving devices - PDAs, laptops, or ambulance-based terminals

Data can be displayed in real time and integrated into the developing pre-hospital patient care record

Can be programmed to process the vital sign data (and provide alerts)

CodeBlue Network Architecture

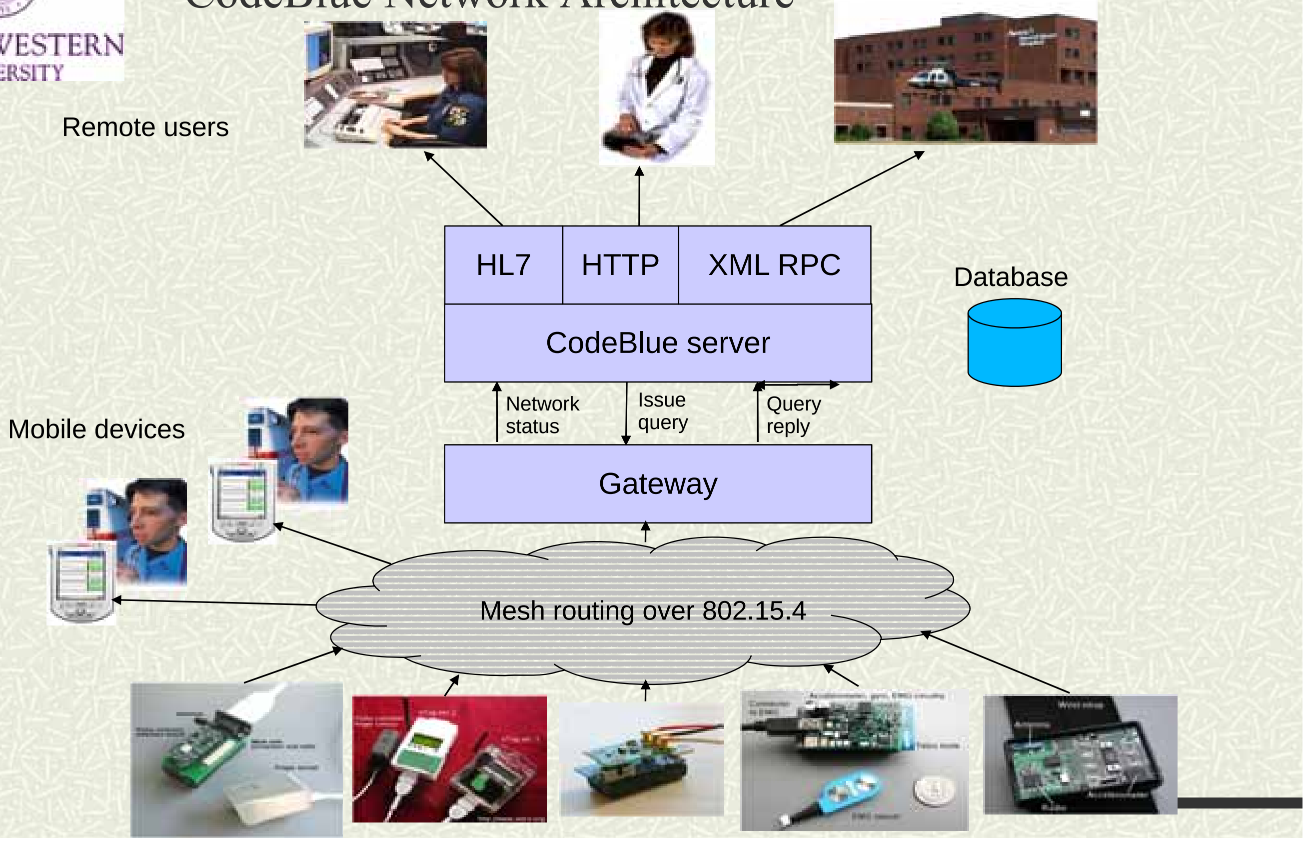

IWESTERN
IVERSITY

What makes this difficult?

- Device mobility
  - Both patient sensors and receiving devices are moving around
  - Need to maintain good connectivity – in elevators, stairwells, etc.
- Multihop, multicast communications
  - Cannot always assume fixed infrastructure – e.g., disaster response
  - Multiple patient sensors may be monitored by multiple end users
- Limited device capabilities
  - Low CPU power (10 MIPS), tiny memories (10 KB of RAM)
  - Cannot run sophisticated protocols with lots of state
- Radio bandwidth is very limited
  - Low-power, 802.15.4 radios – max PHY rate of 250 Kbps
  - Drops to 100 Kbps when taking MAC overhead and framing into account
  - A small number of sensors can rapidly saturate the channel

IWESTERN
IVERSITY

Applications: Emergency medicine

- ual tracking of patient status is difficult
  - Current systems are paper, phone, radio based
  - No real-time updates on patient condition
- ge accidents, fires, terrorist attacks
  - Normal organized community support may be damaged or destroyed
  - Large numbers of patients, severe load on emergency personnel

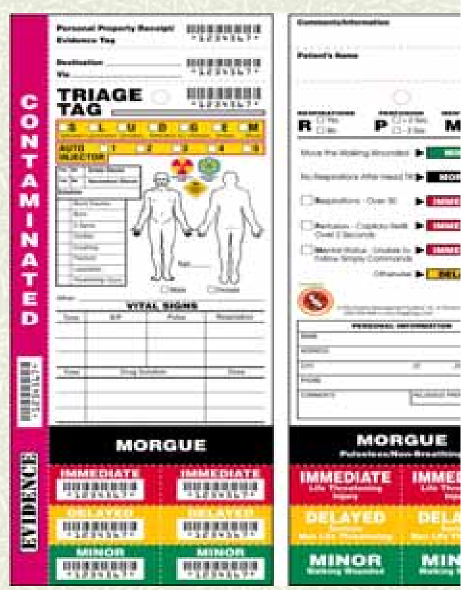

Parkinson's Disease and Stroke Rehab Monitoring

with P. Bonato, Spaulding Rehabilitation Hospital

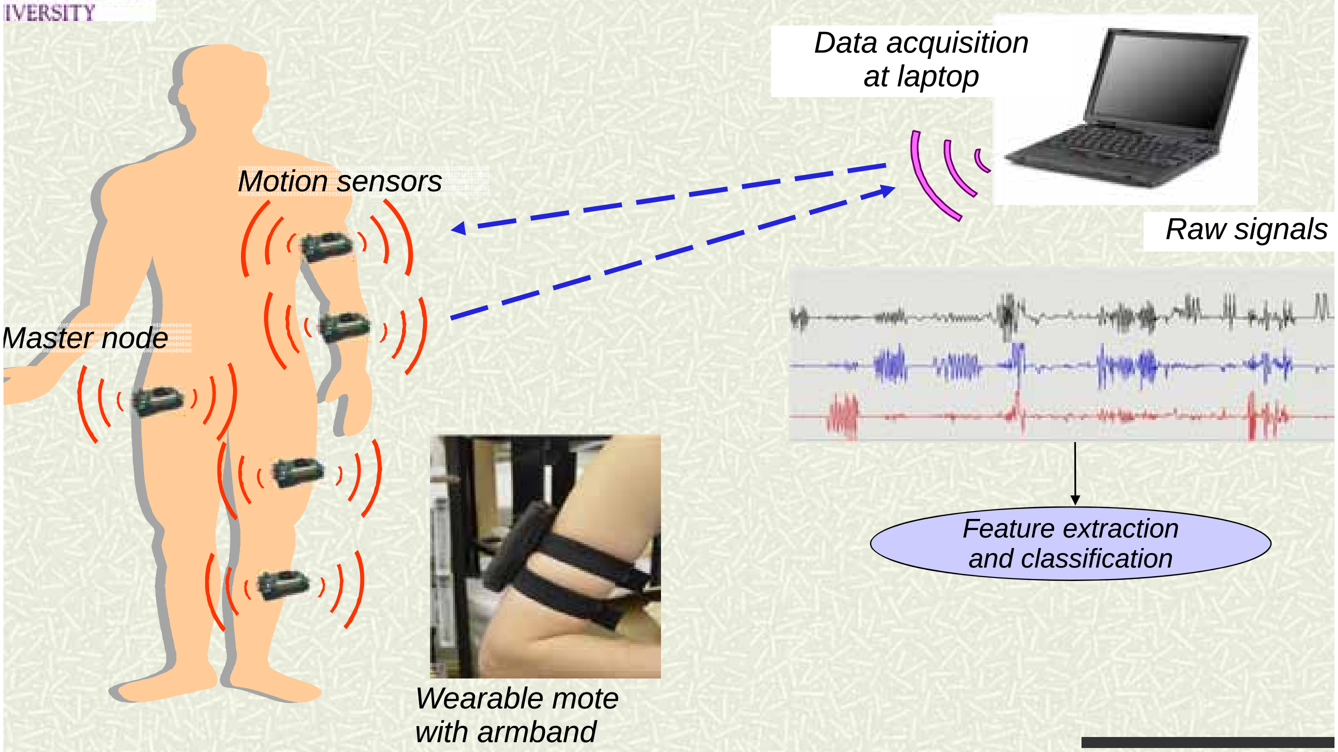

Routing Protocol Design

- CodeBlue requires an *ad hoc* multicast routing protocol
  - *Ad hoc*: No need for fixed infrastructure, forms routes "on demand"
  - Multicast: Data from each sensor can be received by multiple end-user devices
- Ad hoc routing has been extensively studied in wireless environments
  - AODV, CSR, DSDV, ODMRP, ADMR, ...
  - Much of this work done in simulation assuming perfect radio links
  - Implementations primarily focus on laptops or PDAs with 802.11 radios
- What's new here?
  - Very limited radio bandwidth: protocol overhead is a big deal
  - Real radios with lossy, asymmetric links
  - Nodes have very small memory (< 10KB) and limited computational power

The MoteTrack location system

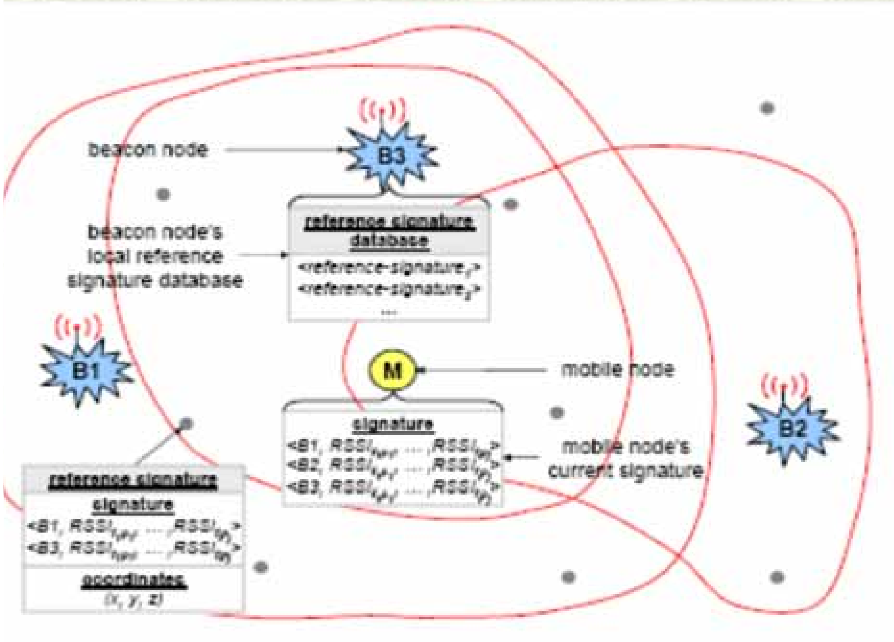

- B1, B2 and B3 are beacons nodes which broadcast beacon messages
- Mobile node M aggregates beacon messages into a signature.
- Comparing the signature and reference signatures stored in beacons nodes, the location of mobile nodes can be determined

The Pothole Patrol: Using a Mobile Sensor Network for Road Surface Monitoring

Hazardous to drivers and increasing repair costs due to vehicle damage

Determine “which” roads need to be fixed

Static sensors will not do well – requires mobility!

P^2 is first of its kind

Challenge : differentiate potholes from other road anomalies *(railroad crossings, expansion joints)*

Challenge : coping with variations in detecting the same pothole. *(speed, sensor orientation)*

P^2 successfully detects most potholes *(>90% accuracy on test data)*

P^2 Architecture

- Vehicles have GPS and 3-axis accelerometer
  - `<time,location,speed,heading,3-axis acceleration>`
- Opportunistic WiFi/Cellular connections with *dPipe* to cope with netw outag
- Taxi Testbed
  - 7 Toyota Priuses[1]
  - Soekris 4801[2] Embedded Linux
  - Wifi Card
  - Sprint EVDO Rev A[3] Network card
  - GPS
- Some numerical facts
  - 9730 total kms
  - 2492 distinct kms
  - 7 cabs
  - 174 km with >10 repeated passes

1. http://www.carbuyersnotebook.com/archives/Toyota_Pruis_2006.jpg
2. http://www.pkgbox.org/Soekris-4801.jpg
3. http://gizmodo.com/gadgets/peripherals/two-new-sprint-evdo-rev-a-cards-pantech-px500-and-sierra-wireless-aircard-595-200423.php

P^2 Architecture

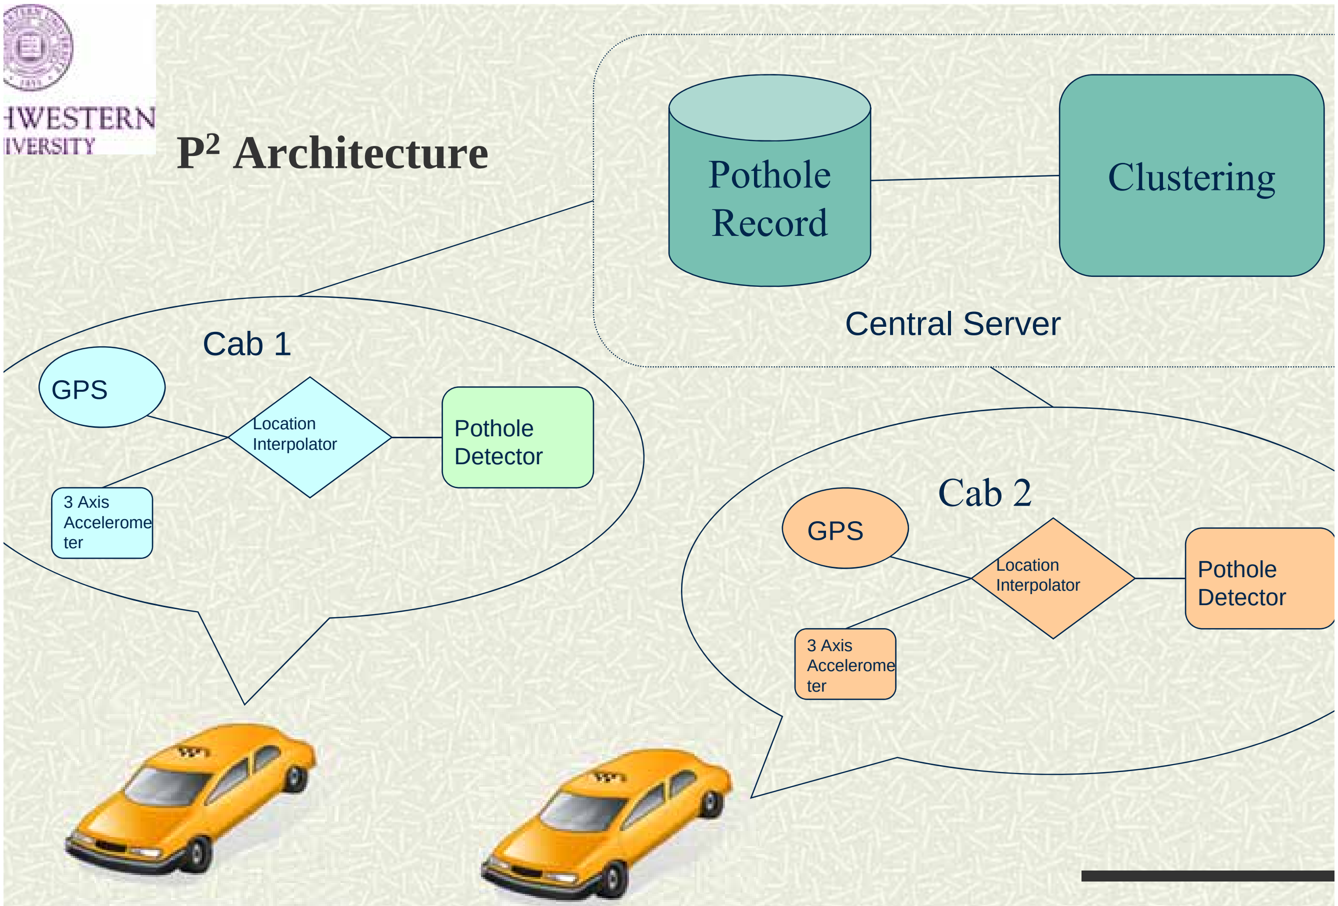

DATA ACQUISITION

- Accelerometer placement
  - Dashboard
  - Windshield
  - Embedded Computer
- GPS Accuracy
  - Standard deviation 3.3m

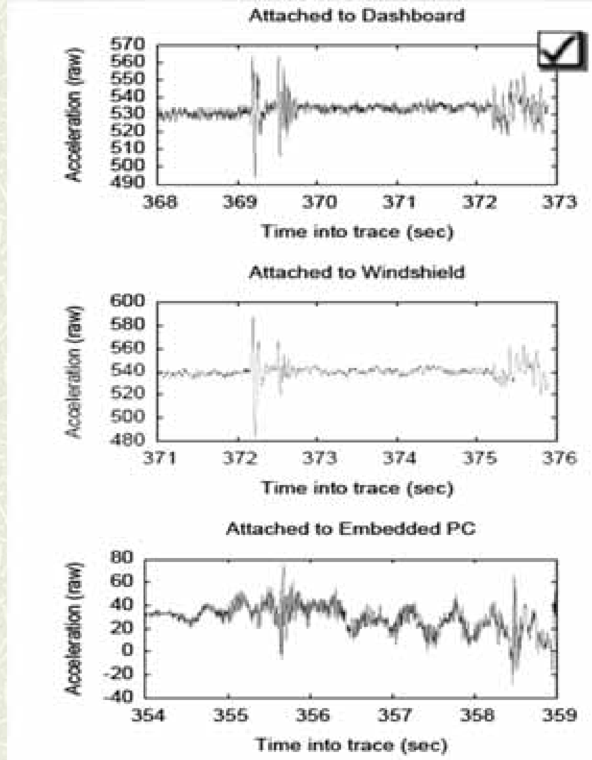

DATA ACQUISITION

- Hand Labeled Data
  - Smooth Road
  - Crosswalks/Expansion Joints
  - Railroad crossing
  - Potholes
  - Manholes
  - Hard Stop
  - Turn

Type	Count	Percentage
Smooth road (SM)	64	23%
Potholes (PH)	63	23%
Manholes (MH)	59	21%
Railroad Crossing (RC)	18	6%
Crosswalk/Exp. Joint (CWEJ)	76	27%

Table 1: Distribution of hand-labeled training data.

DATA ACQUISITION

Loosely Labeled Training Data

- We know only *types of anomalies* and their *rough frequencies*
- Exact numbers and locations are *unknown*
- Extends available training set

- **Storrow Dr.** Heavily used four-lane parkway on the Boston side of the Charles River with several bridges, some potholes.
- **Memorial Dr.** Heavily used four-lane parkway on the Cambridge side of the Charles River, good condition.
- **Binney St.** A two-lane street with many sunk-in manholes and sealed cracks, one pothole.
- **Hwy I-93** An 8 lane interstate highway that cuts through the center of Boston in good condition.
- **Beacham St** A heavily trafficked back road in very poor condition.

ALGORITHM

- Features of accelerometer data
- High energy events are potholes?
  - Not all !
  - Rail road crossings, expansion joints, door slamming are high energy events
- Accelerometer data is processed by embedded computer
  - 256-sample windows
  - Pass through 5 different filters

ALGORITHM - Filtering

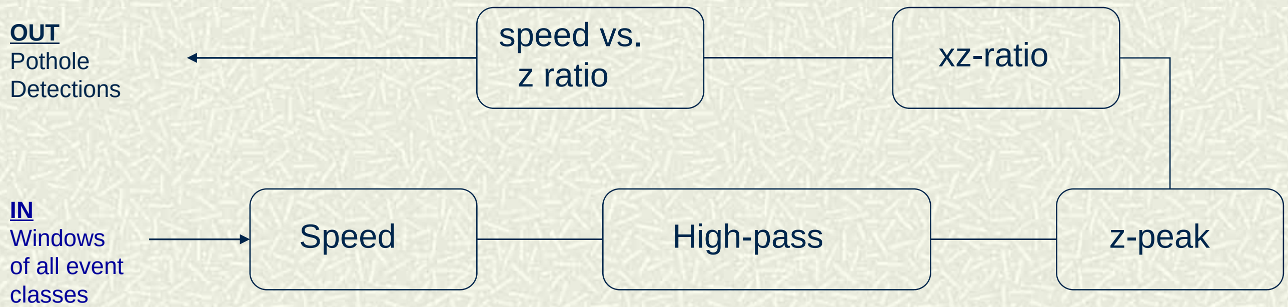

Input

- Raw accelerometer data
- 256-sample windows

❖ Apply signal processing\ filters

ALGORITHM - Filtering

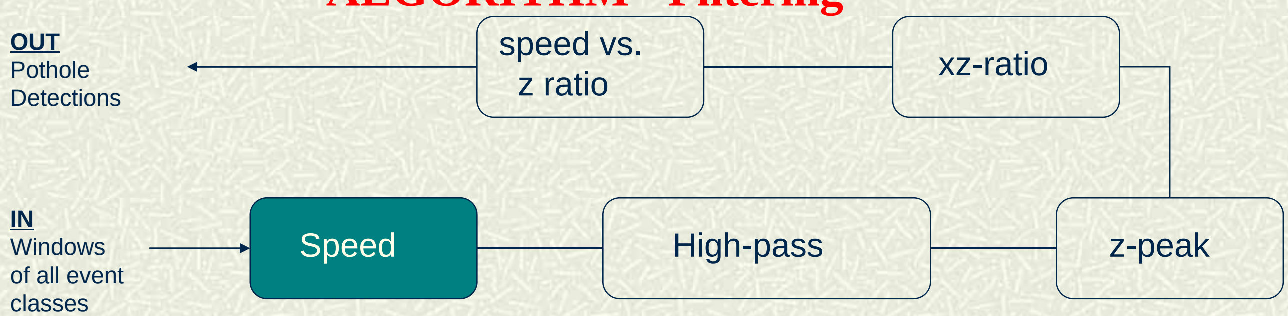

Speed

- Car is not moving or moving slowly
- Rejects door slam and curb ramp events

ALGORITHM - Filtering

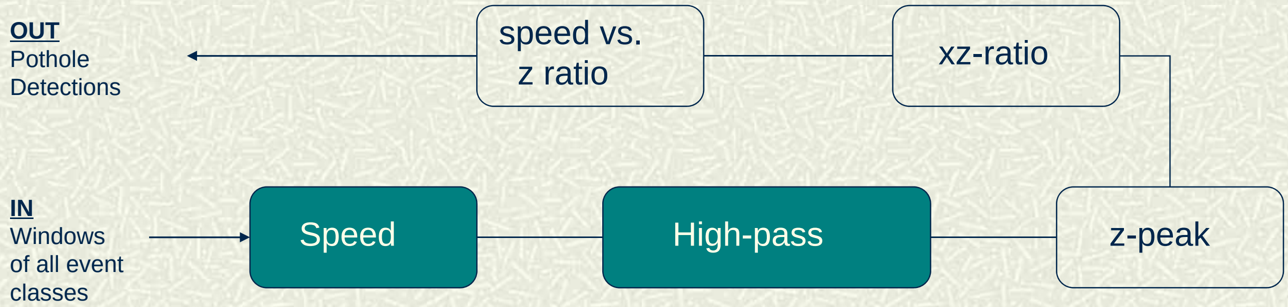

High-Pass

- Removes low-freq components in x and z axes
- Filters out events like turning, veering, braking.

ALGORITHM - Filtering

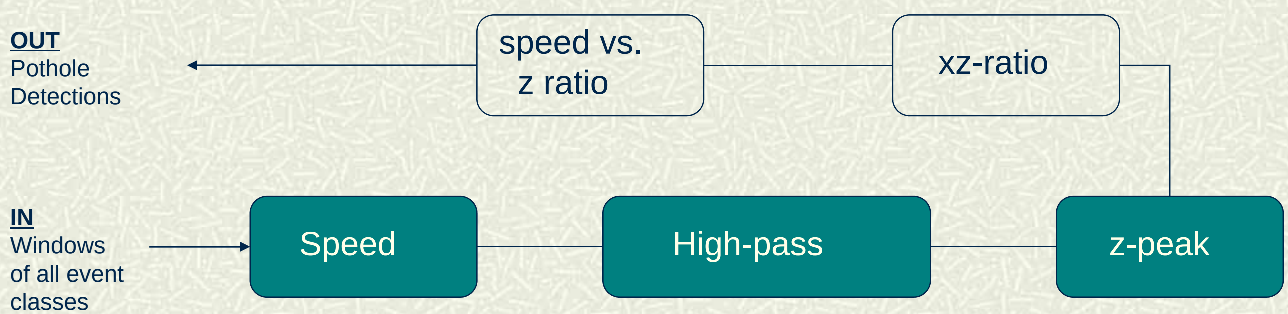

z-peak

- Prime characteristic for significant anomalies
- Rejects all windows with absolute z-acceleration < t_z

ALGORITHM - Filtering

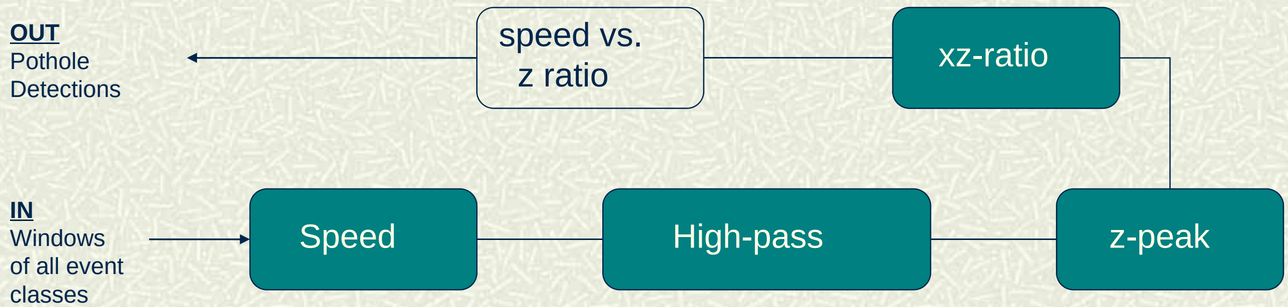

xz- ratio

- Assumes potholes impact only side of the vehicle
- Identifies anomalies that span width of the road (rail crossings, speed bumps)
- Rejects all windows with

  x_{peak} within *Δw (=32)* samples from $z_{peak} < t_x * z_{peak}$

  or, $(X_{peak} / z_{peak}) < t_x$

ALGORITHM - Filtering

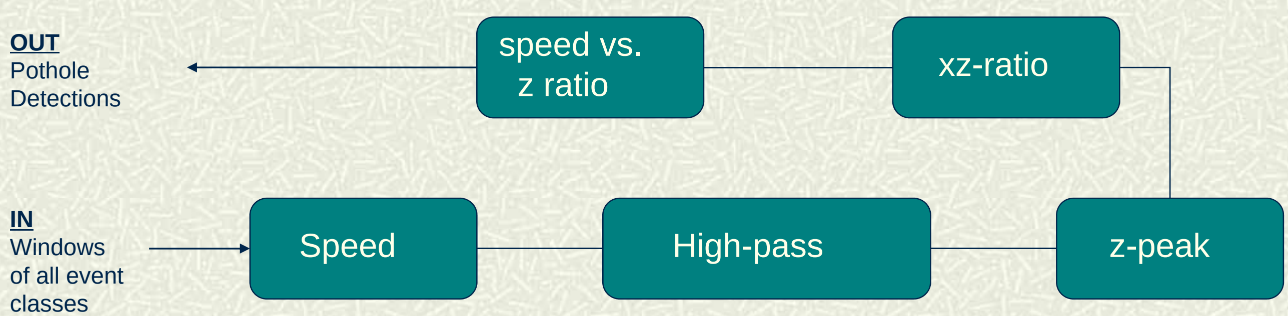

Speed vs. z ratio

- At high speeds, small anomalies cause high peak accelerations
- Rejects windows where $Z_{peak} < t_s * speed$ *or*

  $(Z_{peak} / speed) < t_s$

IWESTERN
IVERSITY

ALGORITHM – Sample Traces

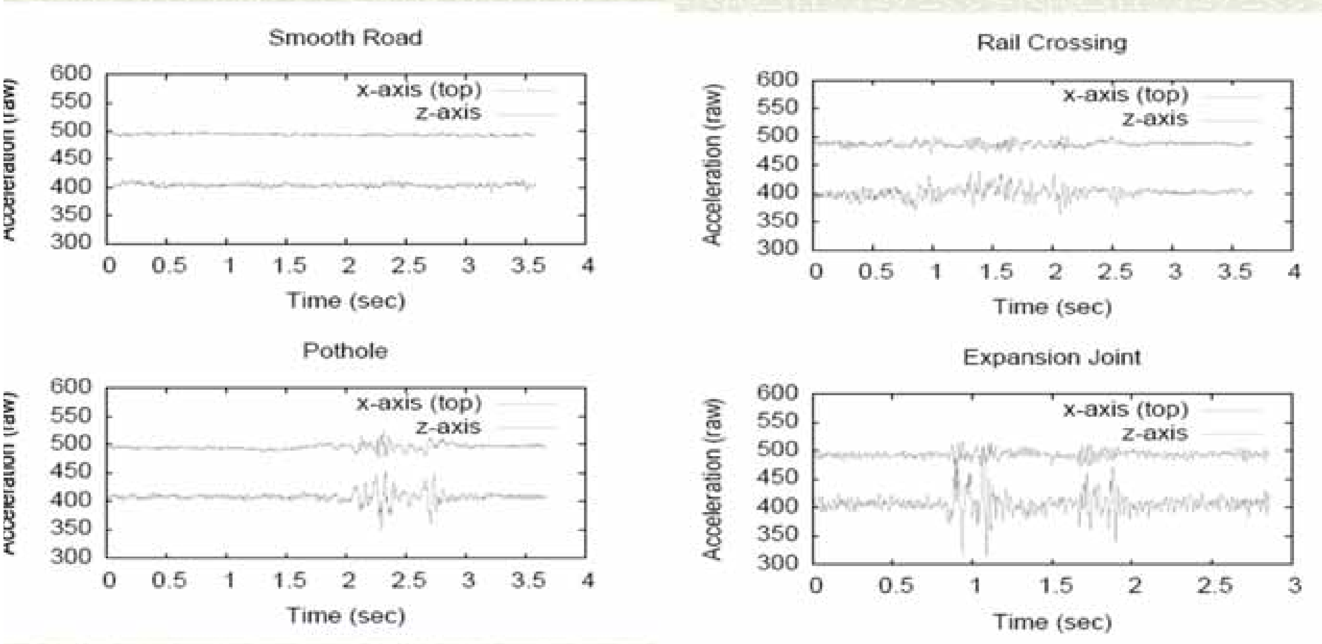

IWESTERN
IVERSITY

ALGORITHM - Training

- Tuning parameters **t**={t_z,t_x,t_s} are computed using exhaustive search over a set of values
- For each set t, we compute detector score
  *s(**t**) = corr – incorr²*
- Corr is no. of pothole detections when sample was labeled as “pothole”
- **Maximize s(t)**
- Include loosely labeled data
  $s(\mathbf{t}) = corr - incorr^{2}_{labeled} - max(0, incorr_{loose} - count_r)$
  where $count_r$ is the expected number of detection in loosely labeled data

ALGORITHM - Clustering

- Improve accuracy
- Cluster of at least *k* events must happen in the same location with small margin of error(Δd)
- Clustering algorithm
  - Place each detection in Δd *X* Δd grid.
  - Compute pairwise distances in same or neighboring grid cells
  - Iteratively merge pairs of distances in order of distance
  - Max intra cluster distance < Δt
  - Reported location is the centroid of the locations within it

Major topics and design principles

- Self-configuration:
  --localization, synchronization, coverage
- In-network programming
  --data aggregation, fusion
- Low-energy data routing
  --data centric approach, geographical routing

Fine Grained Time and Location

- Unlike Internet, node time/space location essential for local/collaborative detection
  - fine-grained localization and time synchronization needed to detect events in three space and compare detections across nodes
- GPS provides solution where available (with differential GPS providing finer granularity)
  - GPS not always available, too "costly," too bulky
  - other approaches under study
- **Localization of sensor nodes has many uses**
  - beamforming for localization of targets and events
  - geographical forwarding
  - geographical addressing

Coverage measures

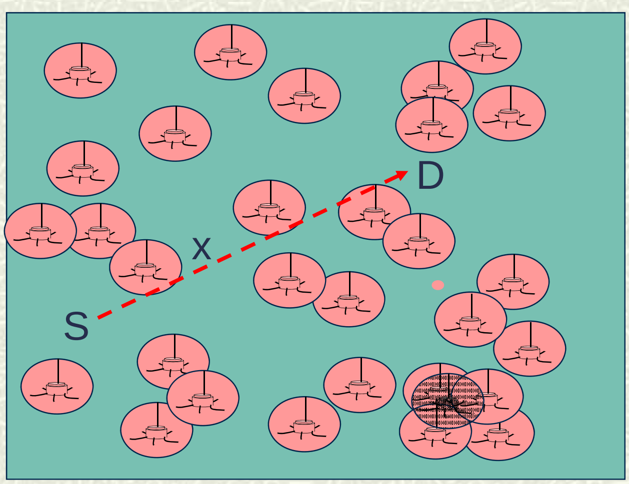

Given: sensor field (either known sensor locations, or spatial density)

- *area coverage:* fraction of area covered by sensors
- detectability: probability sensors detect moving objects
- node coverage: fraction of sensors covered by other sensors
- control:
  - where to add new nodes for max coverage
  - how to move existing nodes for max coverage

In Network Processing

- communication expensive when limited
  - power
  - bandwidth
- perform (data) processing in network
  - close to (at) data
  - forward fused/synthesized results
  - e.g., find max. of data
- distributed data, distributed computation

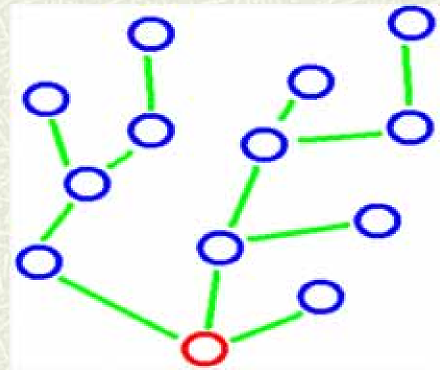

Distributed Representation and Storage

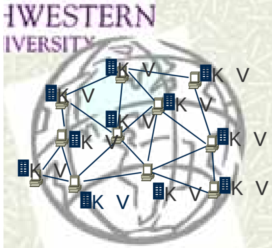

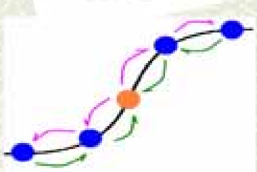

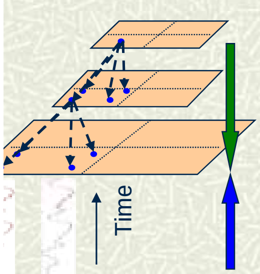

- **Data Centric Protocols, In-network Processing goal:**
  - Interpretation of spatially distributed data (Per-node processing alone is not enough)
  - network does in-network processing based on distribution of data
  - Queries automatically directed towards nodes that maintain relevant/matching data
- **Pattern-triggered data collection**
  - Multi-resolution data storage and retrieval
  - Distributed edge/feature detection
  - Index data for easy temporal and spatial searching
  - Finding global statistics (e.g., distribution)

Directed Diffusion: Data Centric Routing

- **Basic idea**
  - name data (not nodes) with externally relevant attributes: data type, time, location of node, SNR,
  - diffuse requests and responses across network using application driven routing (e.g., geo sensitive or not)
  - support in-network aggregation and processing

- **Data sources publish data, data clients subscribe to data**
  - however, all nodes may play both roles
    - node that aggregates/combines/processes incoming sensor node data becomes a source of new data
    - node that only publishes when combination of conditions arise, is client for triggering event data
  - t*rue peer to peer system?*